Bayesian Semiparametric Multivariate Density Deconvolution

Abhra Sarkar
Department of Statistics, Texas A&M University, 3143 TAMU, College Station, TX 77843-3143 USA
abhra@stat.tamu.edu

Debdeep Pati
Department of Statistics, Florida State University, Tallahassee, FL 32306-4330 USA
debdeep@stat.fsu.edu

Bani K. Mallick
Department of Statistics, Texas A&M University, 3143 TAMU, College Station, TX 77843-3143 USA
bmallick@stat.tamu.edu

Raymond J. Carroll
Department of Statistics, Texas A&M University, 3143 TAMU, College Station, TX 77843-3143 USA
carroll@stat.tamu.edu

arXiv:1404.6462v3 [stat.ME] 16 Nov 2015

Abstract

We consider the problem of multivariate density deconvolution when interest lies in estimating the distribution of a vector valued random variable $\mathbf{X}$ but precise measurements on $\mathbf{X}$ are not available, observations being contaminated by measurement errors $\mathbf{U}$. The existing sparse literature on the problem assumes the density of the measurement errors to be completely known. We propose robust Bayesian semiparametric multivariate deconvolution approaches when the measurement error density of $\mathbf{U}$ is not known but replicated proxies are available for at least some individuals. Additionally, we allow the variability of $\mathbf{U}$ to depend on the associated unobserved values of $\mathbf{X}$ through unknown relationships. Basic properties of finite mixture models, multivariate normal kernels and exchangeable priors are exploited in novel ways to meet modeling and computational challenges. Theoretical results showing the flexibility of the proposed methods in capturing a wide variety of data generating processes are provided. We illustrate the efficiency of the proposed methods in recovering the density of $\mathbf{X}$ through simulation experiments. The methodology is applied to estimate the joint consumption pattern of different dietary components from contaminated 24 hour recalls.

Some Key Words: B-splines, Conditional heteroscedasticity, Latent factor analyzers, Measurement errors, Mixture models, Multivariate density deconvolution, Regularization, Shrinkage.

Short Title: Multivariate Density Deconvolution

1 Introduction

Many problems of practical importance require estimation of the density $f_{\mathbf{X}}$ of a vector valued random variable $\mathbf{X}$. Precise measurements on $\mathbf{X}$ may not, however, be available, observations being contaminated by measurement errors $\mathbf{U}$. Under the assumption of additive measurement errors, the observations are generated from a convolution of the density $f_{\mathbf{X}}$ of $\mathbf{X}$ and the density $f_{\mathbf{U}}$ of the measurement errors $\mathbf{U}$. The problem of estimating the density $f_{\mathbf{X}}$ from available contaminated measurements then becomes a problem of multivariate density deconvolution.

This article proposes novel Bayesian semiparametric density deconvolution approaches based on finite mixtures of latent factor analyzers for robust estimation of the density $f_{\mathbf{X}}$ when the measurement error density $f_{\mathbf{U}}$ is not known but replicated proxies contaminated with measurement errors $\mathbf{U}$ are available for at least some individuals. The proposed deconvolution approaches are highly robust, not having to impose restrictive parametric assumptions on $f_{\mathbf{X}}$ or $f_{\mathbf{U}}$. Additionally, the variability of $\mathbf{U}$ is allowed to depend on the associated unobserved values of $\mathbf{X}$ through unknown relationships. To the best of our knowledge, all existing multivariate deconvolution approaches assume that $\mathbf{U}$ is independent of $\mathbf{X}$ and that the error density $f_{\mathbf{U}}$ is completely known. Ours is thus the first paper that allows the density of the measurement errors to be unknown and free from parametric laws and additionally also accommodates conditional heteroscedasticity in the measurement errors.

The literature on the problem of univariate density deconvolution, in which context we denote the variable of interest by X and the measurement errors by U, is vast. Most of the early literature considered scenarios when the measurement error density f_U is completely known. Fourier inversion based deconvoluting kernel density estimators have been studied by Carroll and Hall (1988), Liu and Taylor (1989), Devroye (1989), Fan (1991a, 1991b, 1992) and Hesse (1999) among many others. For a review of these methods, the reader may be referred to Section 12.1 in Carroll, et al. (2006) and Section 10.2.3 in Buonaccorsi (2010). In reality f_U is rarely known. The problem of deconvolution when the errors are homoscedastic with an unknown density and replicated proxies are available for each subject has been addressed by Li and Vuong (1998). See also Diggle and Hall (1993), Neumann (1997), Carroll and Hall (2004) and the references therein. The assumptions of homoscedasticity of U and their independence from X are also often unrealistic. Flexible Bayesian density deconvolution approaches that allow U to be conditionally heteroscedastic have recently been developed in Staudenmayer, et al. (2008) and Sarkar, et al. (2014). Staudenmayer, et al. (2008) assumed the measurement errors to be normally distributed and used finite mixtures of B-splines to estimate f_X and a variance function that captured the conditional heteroscedasticity. Sarkar, et al. (2014) further relaxed the assumption of normality of U employing flexible infinite mixtures of normal kernels induced by Dirichlet processes to estimate both f_X and f_U.

In sharp contrast to the univariate case, the literature on multivariate density deconvolution is quite sparse. We can only mention Masry (1991), Youndjé and Wells (2008), Comte and Lacour (2013) and Bovy, et al. (2011). The first three considered deconvoluting kernel based approaches assuming the measurement errors $\mathbf{U}$ to be distributed independently from $\mathbf{X}$ according to a completely known probability law. Bovy, et al. (2011) modeled the density $f_{\mathbf{X}}$ using flexible mixtures of multivariate normal kernels, but they assumed $f_{\mathbf{U}}$ to be multivariate normal with known covariance matrices, independent from $\mathbf{X}$. As in the case of univariate problems, the assumptions of a fully specified $f_{\mathbf{U}}$, known covariance matrices, and independence from $\mathbf{X}$ are highly restrictive for most practical applications.

The focus of this article is on multivariate density deconvolution when $f_{\mathbf{U}}$ is not known but replicated proxies are available for at least some individuals. The proposed deconvolution approaches can additionally accommodate conditional heteroscedasticity in $\mathbf{U}$. The problem is important, for instance, in nutritional epidemiology, where nutritionists are typically interested not just in the consumption behaviors of individual dietary components but also in their joint consumption patterns. The data are usually available in the form of dietary recalls and are contaminated by measurement errors that show strong patterns of conditional heteroscedasticity.

As in Sarkar, et al. (2014), we use mixture models to estimate both $f_{\mathbf{X}}$ and $f_{\mathbf{U}}$ but the multivariate nature of the problem brings in new modeling challenges and computational obstacles that preclude straightforward extension of their univariate deconvolution approaches. Instead of using infinite mixtures induced by Dirichlet processes, we use finite mixtures of multivariate normal kernels with exchangeable Dirichlet priors on the mixture probabilities. The use of finite mixtures and exchangeable priors greatly reduces computational complexity while retaining essentially the same flexibility as that of their infinite dimensional counterparts. Carefully constructed priors also allow us to achieve automatic model selection and model averaging. To save space, detailed discussions on these subtle but important issues are moved to Section S.1 in the Supplementary Materials. Additionally, we also exploit the exchangeability of symmetric Dirichlet priors and some basic properties of multivariate normal distributions and finite mixture models to develop a novel strategy that enables us to enforce a required zero mean restriction on the measurement errors. Our proposed technique, as opposed to the one adopted by Sarkar, et al. (2014), is particularly suitable for high dimensional applications and can be easily generalized to enforce moment restrictions of other types on finite mixture models.

It is well known that inverse Wishart priors, due to their dense parametrization, are not suitable for modeling covariance matrices in high dimensional applications. In deconvolution problems the issue is further complicated since $\mathbf{X}$ and $\mathbf{U}$ are both latent. This results in numerically unstable estimates even for small dimensions, particularly when the true covariance matrices are sparse and the likelihood function is of complicated form. To reduce

the effective number of parameters required to be estimated, we consider factor-analytic representation of the component specific covariance matrices with sparsity inducing shrinkage priors on the factor loading matrices.

When we accommodate conditional heteroscedasticity in $\mathbf{U}$, some characteristics of $\mathbf{U}$ prohibit us from using covariance regression techniques. The models that we formulate cater to these properties but result in a complicated form for the likelihood function that again gives rise to significant computational hurdles. We overcome these obstacles by designing a novel two-stage procedure that exploits the same unique properties of conditionally heteroscedastic multivariate measurement errors to our advantage. The procedure first estimates some components of the model using reparametrized versions of the corresponding univariate submodels. The estimates obtained in the first stage are then plugged-in to estimate the remaining parameters in the second stage. Having two separate estimation stages, in some sense, our deconvolution models for conditionally heteroscedastic measurement errors are not purely Bayesian. But they show good empirical performance and, with no other solution available in the existing literature, they provide at least workable starting points towards more sophisticated methodology.

The article is organized as follows. Section 2 details the models. Section 3 presents theoretical results showing flexibility of the proposed models. Section 4 discusses the choice of hyper-parameters. In Section 5, we describe Markov chain Monte Carlo (MCMC) algorithms to sample from the posterior. In Section 6, we detail the two stage estimation procedure for conditionally heteroscedastic measurement errors. Simulation studies comparing the proposed deconvolution methods to a naive method that ignores measurement errors are presented in Section 7. Section 8 presents an application of the proposed methodology in estimation of the joint consumption pattern of dietary intakes from contaminated 24 hour recalls in a nutritional epidemiologic study. Section 9 includes a discussion. An unnumbered section concludes the article with a description of the Supplementary Materials.

2 Deconvolution Models

The goal is to estimate the unknown joint density of a p-dimensional multivariate random variable $\mathbf{X}$. There are $i = 1, \ldots, n$ subjects. Precise measurements of $\mathbf{X}$ are not available. Instead, for $j = 1, \ldots, m_i$, replicated proxies $\mathbf{W}_{ij}$ contaminated with measurement errors $\mathbf{U}_{ij}$ are available for each subject. The replicates are assumed to be generated by the model

$$\mathbf{W}_{ij} \;=\; \mathbf{X}_i + \mathbf{U}_{ij}. \tag{1}$$

Given $\mathbf{X}_i$, $\mathbf{U}_{ij}$ are independently distributed with $E(\mathbf{U}_{ij} \mid \mathbf{X}_i) = \mathbf{0}$. The marginal density of $\mathbf{W}_{ij}$ is denoted by $f_{\mathbf{W}}$. The implied conditional distributions of $\mathbf{W}_{ij}$ and $\mathbf{U}_{ij}$, given $\mathbf{X}_i$, are

denoted by $f_{\mathbf{W}|\mathbf{X}}$ and $f_{\mathbf{U}|\mathbf{X}}$, respectively.

2.1 Modeling the Density $f_{\mathbf{X}}$

In this article $f_{\mathbf{X}}$ is specified as a mixture of multivariate normal kernels

$$f_{\mathbf{X}}(\mathbf{X}) = \sum_{k=1}^{K_{\mathbf{X}}} \pi_{\mathbf{X},k} \ \mathrm{MVN}_p(\mathbf{X} \mid \boldsymbol{\mu}_{\mathbf{X},k}, \boldsymbol{\Sigma}_{\mathbf{X},k}), \tag{2}$$

where $\mathrm{MVN}_p(\cdot \mid \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes a p-dimensional multivariate normal density with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. For the rest of this subsection, the subscript $\mathbf{X}$ is kept implicit to keep the notation clean.

We assign a finite Dirichlet prior to the mixture probability vector $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_K)^{\mathrm{T}}$ as

$$\boldsymbol{\pi} \ \sim \ \mathrm{Dir}(\alpha/K, \ldots, \alpha/K). \tag{3}$$

Here $\mathrm{Dir}(\alpha_1, \ldots, \alpha_K)$ denotes a finite dimensional Dirichlet distribution on the K-dimensional unit simplex with concentration parameter $(\alpha_1, \ldots, \alpha_K)$. Given K and the latent cluster membership indices, the prior is conjugate. The symmetry of the assumed Dirichlet prior helps in additional reduction of computational complexity by simplifying MCMC mixing issues. Provided K is sufficiently large, a carefully chosen α can impart the posterior with certain properties that simplify model selection and model averaging issues by influencing the posterior to concentrate in regions that favor empty redundant components, see Section 4 and Section S.1 of the Supplementary Materials. We assign conjugate multivariate normal priors to the component specific mean vectors $\boldsymbol{\mu}_k$, so that

$$\boldsymbol{\mu}_k \ \sim \ \mathrm{MVN}_p(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0). \tag{4}$$

The conjugacy again helps in simplifying posterior calculations. Later on, we will employ similar mixture models for the density of the measurement errors, and this conjugacy, along with some basic properties of multivariate normal kernels, will also help us enforce the mean zero restriction on the measurement errors. For the component specific covariance matrices $\boldsymbol{\Sigma}_k$, we first consider conjugate inverse Wishart priors

$$\boldsymbol{\Sigma}_k \sim \mathrm{IW}_p(\nu_0, \boldsymbol{\Psi}_0). \tag{5}$$

Here $\mathrm{IW}_p(\nu, \boldsymbol{\Psi})$ denotes an inverse Wishart density on the space of $p \times p$ positive definite matrices with mean $\boldsymbol{\Psi}/(\nu - p - 1)$. While the conjugacy of the inverse Wishart priors helps in simplifying posterior calculations, in complex high dimensional problems its dense parameterization may result in numerically unstable estimates, particularly when the covariance matrices are sparse. In a deconvolution problem the issue is compounded further by the nonavailability of the true $\mathbf{X}_i$'s. To reduce the effective number of parameters to be estimated, we consider a parsimonious factor-analytic representation of the component specific

covariance matrices:

$$\boldsymbol{\Sigma}_k = \boldsymbol{\Lambda}_k \boldsymbol{\Lambda}_k^{\mathrm{T}} + \boldsymbol{\Omega}, \tag{6}$$

where $\boldsymbol{\Lambda}_k$ are $p \times q_k$ factor loading matrices and $\boldsymbol{\Omega}$ is a diagonal matrix with non-negative entries. In practical applications q_k will typically be much smaller than p, inducing parsimonious characterizations of the unknown covariance matrices $\boldsymbol{\Sigma}_k$. Model (2) can be equivalently represented as

$$\Pr(C_i = k) = \pi_k, \tag{7}$$

$$(\mathbf{X}_i \mid C_i = k) = \boldsymbol{\mu}_k + \boldsymbol{\Lambda}_k \boldsymbol{\eta}_i + \boldsymbol{\Delta}_i, \tag{8}$$

$$\boldsymbol{\eta}_i \sim \mathrm{MVN}_p(\mathbf{0}, \mathbf{I}_p), \qquad \boldsymbol{\Delta}_i \sim \mathrm{MVN}_p(\mathbf{0}, \boldsymbol{\Omega}), \tag{9}$$

where C_i are the mixture labels associated with $\mathbf{X}_i$, $\boldsymbol{\eta}_i$ are latent factors, and $\boldsymbol{\Delta}_i$ are errors with covariance $\boldsymbol{\Omega} = \mathrm{diag}(\sigma_1^2, \ldots, \sigma_p^2)$.

The above characterization of $\boldsymbol{\Sigma}_k$ is not unique, since for any semi-orthogonal matrix $\mathbf{P}$ the loading matrix $\boldsymbol{\Lambda}_k^1 = \boldsymbol{\Lambda}_k \mathbf{P}$ also satisfies (6). However, in a Bayesian framework, when interest lies primarily in estimating the density $f_{\mathbf{X}}$, identifiability of the latent factors is not required. This also allows the loading matrices to have a-priori a potentially infinite number of columns. Sparsity inducing priors, that favor more shrinkage as the column index increases, can then be used to shrink the redundant columns towards zero. In this article, we do this by adapting the shrinkage priors proposed in Bhattacharya and Dunson (2011) that allow easy posterior computation. Let $\boldsymbol{\Lambda}_k = ((\lambda_{k,jh}))_{j=1,h=1}^{p,\infty}$, where j and h denote the row and the column indices, respectively. For $h = 1, \ldots, \infty$, we assign priors as follows

$$\lambda_{k,jh} \sim \mathrm{Normal}(0, \phi_{k,jh}^{-1} \tau_{k,h}^{-1}), \qquad \phi_{k,jh} \sim \mathrm{Ga}(\nu/2, \nu/2), \tag{10}$$

$$\tau_{k,h} \sim \prod_{\ell=1}^{h} \delta_{k,\ell}, \qquad \delta_{k,\ell} \sim \mathrm{Ga}(a_\ell, 1), \qquad \sigma_j^2 \sim \text{Inv-Ga}(a_\sigma, b_\sigma). \tag{11}$$

Here $\mathrm{Ga}(\alpha, \beta)$ denotes a Gamma distribution with shape parameter α and rate parameter β and $\mathrm{IG}(a, b)$ denotes an inverse-Gamma distribution with shape parameter a and scale parameter b. In the k^{th} component factor loading matrix $\boldsymbol{\Lambda}_k$, the parameters $\{\phi_{k,jh}\}_{j=1}^{p}$ control the local shrinkage of the elements in the h^{th} column, whereas $\tau_{k,h}$ controls the global shrinkage. When $a_h > 1$ for $h = 2, \ldots, \infty$, the sequence $\{\tau_{k,h}\}_{h=1}^{\infty}$ becomes stochastically increasing and thus favors more shrinkage as the column index h increases.

In addition to inducing adaptive sparsity and hence numerical stability, by favoring more shrinkage as the column index increases, the shrinkage priors play another important role in making the proposed factor analytic model highly robust to misspecification of the number of latent factors, allowing us to adopt simple strategies to determine the number of latent factors to be included in the model in practice. Details are deferred to Section 4.

Throughout the rest of the paper, mixtures with inverse Wishart prior on the covariance matrices will be referred to as MIW models and mixtures of latent factor analyzers will be

referred to as MLFA models.

For a review of finite mixture models and mixtures of latent factor analyzers, without moment restrictions or sparsity inducing priors and with applications in measurement error free scenarios, see Fokoué and Titterington (2003), Frühwirth-Schnatter (2006), Mengersen, et al. (2011) and the references therein. For other types of shrinkage priors, see Brown and Griffin (2010), Carvalho, et al. (2010), Bhattacharya, et al. (2014) etc.

2.2 Modeling the Density $f_{\mathbf{U}}$ of the Measurement Errors

2.2.1 Independently Distributed Measurement Errors

In this section, we develop models for the measurement errors $\mathbf{U}$ assuming them to be independent from $\mathbf{X}$. That is, we assume $f_{\mathbf{U}|\mathbf{X}} = f_{\mathbf{U}}$ for all $\mathbf{X}$. This remains the most extensively researched deconvolution problem for both univariate and multivariate cases. The techniques developed in this section will also provide crucial building blocks for more realistic models in Section 2.2.2. The measurement errors and their density are now denoted by $\boldsymbol{\epsilon}_{ij}$ and $f_{\boldsymbol{\epsilon}}$, respectively, for reasons to become obvious shortly in Section 2.2.2.

As in Section 2.1, a mixture of multivariate normals can be used to model the density $f_{\boldsymbol{\epsilon}}$ but the model now has to satisfy a mean zero constraint. That is

$$f_{\boldsymbol{\epsilon}}(\boldsymbol{\epsilon}) = \textstyle\sum_{k=1}^{K_{\boldsymbol{\epsilon}}} \pi_{\boldsymbol{\epsilon},k}\ \mathrm{MVN}_p(\boldsymbol{\epsilon} \mid \boldsymbol{\mu}_{\boldsymbol{\epsilon},k}, \boldsymbol{\Sigma}_{\boldsymbol{\epsilon},k}), \tag{12}$$

$$\text{subject to } \textstyle\sum_{k=1}^{K_{\boldsymbol{\epsilon}}} \pi_{\boldsymbol{\epsilon},k}\boldsymbol{\mu}_{\boldsymbol{\epsilon},k} = \mathbf{0}. \tag{13}$$

To get numerically stable estimates of the density of the errors, latent factor characterization of the covariance matrices with sparsity inducing shrinkage priors as in Section 2.1 may again be used. Details are curtailed to avoid unnecessary repetition and we only present the mechanism to enforce the zero mean restriction on the model. The subscript $\boldsymbol{\epsilon}$ is again dropped in favor of cleaner notation. In later sections, the subscripts $\mathbf{X}$ and $\boldsymbol{\epsilon}$ reappear to distinguish between the parameters associated with $f_{\mathbf{X}}$ and $f_{\boldsymbol{\epsilon}}$, when necessary.

Without the mean restriction and under conjugate multivariate normal priors $\boldsymbol{\mu}_k \sim \mathrm{MVN}_p(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$, the posterior full conditional of $\boldsymbol{\mu}^{Kp\times 1} = (\boldsymbol{\mu}_1^{\mathrm{T}}, \ldots, \boldsymbol{\mu}_K^{\mathrm{T}})^{\mathrm{T}}$ is given by

$$\mathrm{MVN}_{Kp}\left\{\begin{pmatrix} \boldsymbol{\mu}_1^0 \\ \boldsymbol{\mu}_2^0 \\ \vdots \\ \boldsymbol{\mu}_K^0 \end{pmatrix}, \begin{pmatrix} \boldsymbol{\Sigma}_1^0 & \mathbf{0} & \ldots & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Sigma}_2^0 & \ldots & \mathbf{0} \\ \vdots & & & \\ \mathbf{0} & \mathbf{0} & \ldots & \boldsymbol{\Sigma}_K^0 \end{pmatrix}\right\} \equiv \mathrm{MVN}_{Kp}(\boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0), \tag{14}$$

where $\boldsymbol{\epsilon}_{ij}$ and other conditioning variables are implicitly understood. Explicit expressions of $\boldsymbol{\mu}^0$ and $\boldsymbol{\Sigma}^0$ in terms of the conditioning variables can be found in Section 4. The posterior

full conditional of $\boldsymbol{\mu}$ under the mean restriction can then be obtained easily by further conditioning the distribution in (14) by $\boldsymbol{\mu}_R = \sum_{k=1}^{K} \pi_k \boldsymbol{\mu}_k = 0$ and is given by

$$(\boldsymbol{\mu} \mid \boldsymbol{\mu}_R = \mathbf{0}) \sim \mathrm{MVN}_{Kp}\{\boldsymbol{\mu}^0 - \boldsymbol{\Sigma}^0_{1,R}(\boldsymbol{\Sigma}^0_{R,R})^{-1}\boldsymbol{\mu}^0_R, \boldsymbol{\Sigma}^0 - \boldsymbol{\Sigma}^0_{1,R}(\boldsymbol{\Sigma}^0_{R,R})^{-1}\boldsymbol{\Sigma}^0_{R,1}\}, \tag{15}$$

where $\boldsymbol{\mu}^0_R = \sum_{k=1}^{K} \pi_k \boldsymbol{\mu}^0_k = E(\boldsymbol{\mu}_R)$, $\boldsymbol{\Sigma}_{k,K} = \pi_k \boldsymbol{\Sigma}^0_k = \mathrm{cov}(\boldsymbol{\mu}_k, \boldsymbol{\mu}_R)$, $\boldsymbol{\Sigma}^0_{R,R} = \boldsymbol{\Sigma}_{K+1,K+1} = \sum_{k=1}^{K} \pi_k^2 \boldsymbol{\Sigma}^0_k = \mathrm{cov}(\boldsymbol{\mu}_R)$, and $\boldsymbol{\Sigma}^0_{R,1} = (\boldsymbol{\Sigma}_{1,K+1}, \boldsymbol{\Sigma}_{2,K+1}, \ldots, \boldsymbol{\Sigma}_{K,K+1})$. To sample from this singular density, we can first sample from the non-singular distribution of $\{(\boldsymbol{\mu}_1^{\mathrm{T}}, \boldsymbol{\mu}_2^{\mathrm{T}}, \ldots, \boldsymbol{\mu}_{K-1}^{\mathrm{T}})^{\mathrm{T}} \mid \boldsymbol{\mu}_R = \mathbf{0}\}$, which can also be trivially obtained from (15), and then set $\boldsymbol{\mu}_K = -\sum_{k=1}^{K-1} \pi_k \boldsymbol{\mu}_k / \pi_K$.

Two remarks are in order. First, the symmetric Dirichlet prior on the mixture probabilities plays an additional implicit but important role here. Although we have used the K^{th} component to enforce the mean restriction, under exchangeable Dirichlet priors, the posterior is also invariant to permutations of the mixture labels, making all the components equally deserving candidates for this fix and the specific choice of any particular component irrelevant. Second, the method depends primarily on the properties of the priors on the mixture probabilities and the mean vectors and not on the model for the covariance matrices. The mechanism can therefore be applied quite generally in conjunction with any model for the component specific covariance matrices belonging to a class that does not disturb the label invariance property of the posterior. This class includes both the MIW and the MLFA models described in this article.

2.2.2 Conditionally Heteroscedastic Measurement Errors

We now consider the case when the variances of the measurement errors depend on the associated unknown values of $\mathbf{X}$ through unknown relationships.

Interpreting the conditioning variables $\mathbf{X}$ broadly as predictors, one can loosely connect our problem of modeling conditionally heteroscedastic $\mathbf{U}$ to the problem of covariance regression (Hoff and Niu, 2012; Fox and Dunson, 2013, etc.), where the covariance of the multivariate regression errors are allowed to vary flexibly with precisely measured and possibly multivariate predictors. In covariance regression problems, the dimension of the regression errors is typically unrelated to the dimension of the predictors. In multivariate deconvolution problems, in contrast, the dimension of each $\mathbf{U}_{ij}$ is exactly the same as the dimension of $\mathbf{X}_i$, the ℓ^{th} component $U_{ij\ell}$ being the measurement error associated exclusively with $X_{i\ell}$. This implies that although different components of the measurement error vectors $\mathbf{U}_{ij}$ may be correlated, the dependence of $U_{ij\ell}$ on $\mathbf{X}_i$ can be explained mostly through $X_{i\ell}$. Figure 6, for instance, suggests very strong conditional heteroscedasticity patterns and it is plausible to assume the conditional variability in $U_{ij\ell}$ to be explained mostly through $X_{i\ell}$ only. It is interesting to note these contrasts between conditionally heteroscedastic regression and measurement errors become particularly prominent in the multivariate set up. See Section

S.2 and Section S.3 in the Supplementary Materials for additional discussion.

These differences between conditionally heteroscedastic multivariate regression and measurement errors preclude direct application of covariance regression approaches to model conditionally heteroscedastic measurement errors and require that models that can accommodate their unique characteristics be specially designed. As a plausible model for conditionally heteroscedastic multivariate measurement errors, we assume that

$$(\mathbf{U}_{ij} \mid \mathbf{X}_i) = \mathbf{S}(\mathbf{X}_i)\boldsymbol{\epsilon}_{ij}, \tag{16}$$

where $\mathbf{S}(\mathbf{X}_i) = \text{diag}\{s_1(X_{i1}), s_2(X_{i2}), \ldots, s_p(X_{ip})\}$ and $\boldsymbol{\epsilon}_{ij}$, henceforth referred to as the 'scaled errors', are distributed independently of $\mathbf{X}_i$. Model (16) implies that $\text{cov}(\mathbf{U}_{ij} \mid \mathbf{X}_i) = \mathbf{S}(\mathbf{X}_i)\ \text{cov}(\boldsymbol{\epsilon}_{ij})\ \mathbf{S}(\mathbf{X}_i)$ and marginally $\text{var}(U_{ij\ell} \mid \mathbf{X}_i) = s_\ell^2(X_{i\ell})\text{var}(\epsilon_{ij\ell})$, a function of $X_{i\ell}$ only, as desired. The techniques developed in Section 2.2.1 can now be employed to model the density of $\boldsymbol{\epsilon}_{ij}$, allowing different components of $\mathbf{U}_{ij}$ to be correlated and their joint density to deviate from multivariate normality.

We model the variance functions s_ℓ^2, denoted also by v_ℓ, using positive mixtures of B-spline basis functions with smoothness inducing priors on the coefficients as in Staudenmayer, et al. (2008). For the ℓ^{th} component, partition an interval $[A_\ell, B_\ell]$ of interest into L_ℓ subintervals using knot points $A_\ell = t_{\ell,1} = \cdots = t_{\ell,q+1} < t_{\ell,q+2} < t_{\ell,q+3} < \cdots < t_{\ell,q+L_k} < t_{\ell,q+L_\ell+1} = \cdots = t_{\ell,2q+L_\ell+1} = B_\ell$. A flexible model for the variance functions is given by

$$v_\ell(X_{i\ell}) = s_\ell^2(X_{i\ell}) = \textstyle\sum_{j=1}^{J_\ell} b_{q,j,\ell}(X_{i\ell}) \exp(\xi_{j\ell}) = \mathbf{B}_{q,J_\ell,\ell}(X_{i\ell}) \exp(\boldsymbol{\xi}_\ell), \tag{17}$$

$$(\boldsymbol{\xi}_\ell \mid J_\ell, \sigma_{\xi,\ell}^2) \propto (2\pi\sigma_{\xi,\ell}^2)^{-J_\ell/2} \exp\{-\boldsymbol{\xi}_k^{\mathrm{T}} P_\ell \boldsymbol{\xi}_\ell/(2\sigma_{\xi,\ell}^2)\}, \quad \sigma_{\xi,\ell}^2 \sim \text{Inv-Ga}(a_\xi, b_\xi). \tag{18}$$

Here $\{b_{q,j,\ell}\}_{j=1}^{J_\ell}$ denote $J_\ell = (q + L_\ell)$ B-spline bases of degree q as defined in de Boor (2000), $\boldsymbol{\xi}_\ell = \{\xi_{1\ell}, \xi_{2\ell}, \ldots, \xi_{J_\ell\ell}\}^{\mathrm{T}}$; $\exp(\boldsymbol{\xi}_\ell) = \{\exp(\xi_{1\ell}), \exp(\xi_{2\ell}), \ldots, \exp(\xi_{J_\ell\ell})\}^{\mathrm{T}}$; and $P_\ell = D_\ell^{\mathrm{T}} D_\ell$, where D_ℓ is a $J_\ell \times (J_\ell + 2)$ matrix such that $D_\ell \boldsymbol{\xi}_\ell$ computes the second differences in $\boldsymbol{\xi}_\ell$. The prior $p_0(\boldsymbol{\xi}_\ell \mid \sigma_{\xi,\ell}^2)$ induces smoothness in the coefficients because it penalizes $\sum_{j=1}^{J_k} (\Delta^2 \xi_{j\ell})^2 = \boldsymbol{\xi}_\ell^{\mathrm{T}} P_\ell \boldsymbol{\xi}_\ell$, the sum of squares of the second order differences in $\boldsymbol{\xi}_\ell$ (Eilers and Marx, 1996). The parameters $\sigma_{\xi,\ell}^2$ play the role of smoothing parameter - the smaller the value of $\sigma_{\xi,\ell}^2$, the stronger the penalty and the smoother the variance function. The inverse-Gamma hyperpriors on $\sigma_{\xi,\ell}^2$ allow the data to have influence on the posterior smoothness and make the approach data adaptive.

3 Model Flexibility

This section presents a theoretical study of the flexibility of the proposed models. Proofs of the results are presented in the Supplementary Materials. We focus on the deconvolution models for conditionally heteroscedastic measurement errors, the case of independently dis-

tributed errors following as a special case. First we show that componentwise our models for the density $f_{\mathbf{X}}$ of $\mathbf{X}$, the density $f_{\boldsymbol{\epsilon}}$ of the scaled errors $\boldsymbol{\epsilon}$, and the variance functions v_ℓ are all highly flexible. Building on these results, we then show that our proposed deconvolution models can accommodate a large class of data generating processes.

Let the generic notation Π denote a prior on some class of random functions. Also let $\mathcal{T}$ denote the target class of functions to be modeled by Π. The support of Π throws light on the flexibility of Π. For Π to be a flexible prior, one would expect that $\mathcal{T}$ or a large subset of $\mathcal{T}$ would be contained in the support of Π.

For investigating the flexibility of priors for density functions, a relevant concept is that of Kullback-Leibler (KL) support. The KL divergence between two densities f_0 and f, denoted by $d_{KL}(f_0, f)$, is defined as $d_{KL}(f_0, f) = \int f_0(Z) \log \{f_0(Z)/f(Z)\} dZ$. Let Π_f denote a prior assigned to a random density f. A density f_0 is said to belong to the KL support of Π_f if $\Pi_f\{f : d_{KL}(f_0, f) < \delta\} > 0 \ \forall \delta > 0$. The class of densities in the KL support of Π_f is denoted by $KL(\Pi_f)$.

Let $\mathcal{F}$ be the class of target densities to be modeled by the prior Π_f. Let $\mathcal{S}$ denote the support of $\mathcal{F}$ and $\widetilde{\mathcal{F}} \subseteq \mathcal{F}$ denote the class of densities that satisfy the following fairly minimal set of regularity conditions. Since $\widetilde{\mathcal{F}}$ is a large subclass of $\mathcal{F}$, its inclusion in the KL support of Π_f would establish the flexibility of Π_f.

Conditions 1. *1. f_0 is continuous on $\mathcal{S}$ except on a set of measure zero.*
2. The second order moments of f_0 are finite.
3. For some $r > 0$ and for all $\mathbf{z} \in \mathcal{S}$, there exist hypercubes $C_r(\mathbf{z})$ with side length r and $\mathbf{z} \in C_r(\mathbf{z})$ such that

$$\int f_0(\mathbf{z}) \ log \left\{ \frac{f_0(\mathbf{z})}{\inf_{\mathbf{t} \in C_r(\mathbf{z})} f_0(\mathbf{t})} \right\} d\mathbf{z} < \infty.$$

Let $\Pi_{\mathbf{X}}$ be a generic notation for both the MIW and the MLFA prior on $f_{\mathbf{X}}$ defined in Section 2.1. Similarly, let $\Pi_{\boldsymbol{\epsilon}}$ be a generic notation for both the MIW and the MLFA prior on $f_{\boldsymbol{\epsilon}}$ defined in Section 2.2. When the measurement errors are distributed independently of $\mathbf{X}$, the support of $f_{\mathbf{X}}$, say $\mathcal{X}$, may be taken to be any subset of $\mathbb{R}^p$. For conditionally heteroscedastic measurement errors, the variance functions $s^2_\ell(\cdot)$ that capture the conditional variability are modeled by mixtures of B-splines defined on closed intervals $[A_k, B_k]$. In this case, the support of $f_{\mathbf{X}}$ is assumed to be the closed hypercube $\mathcal{X} = [A_1, B_1] \times \cdots \times [A_p, B_p]$. Let $\mathcal{F}_{\mathbf{X}}$ denote the set of all densities on $\mathcal{X}$, the target class of densities to be modeled by $\Pi_{\mathbf{X}}$ and $\widetilde{\mathcal{F}}_{\mathbf{X}} \subseteq \mathcal{F}_{\mathbf{X}}$ denote the class of densities $f_{0\mathbf{X}}$ that satisfy Conditions 1. Similarly, let $\mathcal{F}_{\boldsymbol{\epsilon}}$ denote the set of all densities on $\mathbb{R}^p$ that have mean zero and $\widetilde{\mathcal{F}}_{\boldsymbol{\epsilon}} \subseteq \mathcal{F}_{\boldsymbol{\epsilon}}$ denote the class of densities $f_{0\boldsymbol{\epsilon}}$ that satisfy Conditions 1. The following Lemma establishes the flexibility of the models for $f_{\mathbf{X}}$ and $f_{\boldsymbol{\epsilon}}$.

Lemma 1. *1. $\widetilde{\mathcal{F}}_{\mathbf{X}} \subseteq KL(\Pi_{\mathbf{X}})$ 2. $\widetilde{\mathcal{F}}_{\boldsymbol{\epsilon}} \subseteq KL(\Pi_{\boldsymbol{\epsilon}})$.*

For investigating the flexibility of models for general classes of functions, a relevant concept is that of sup norm support. The sup norm distance between two functions g_0 and g, denoted by $||g_0 - g||_\infty$, is defined as $||g_0 - g||_\infty = \sup_Z |g_0(Z) - g(Z)|$. Let Π_g denote a prior assigned to a random function g. A function g_0 is said to belong to the sup norm support of Π_g if $\Pi_g(g : ||g_0 - g||_\infty < \delta) > 0 \ \forall \delta > 0$. The class of functions in the sup norm support of Π_g is denoted by $SN(\Pi_g)$.

Let $\Pi_{\mathbf{V}}$ denote the prior on the variance functions based on mixtures of B-spline basis functions defined in Section 2.2.2. For notational convenience we consider the case of a univariate variance function supported on $[A, B]$. Extension to the multivariate case with variance functions supported on $\mathcal{X}$ is technically trivial. Let $\mathcal{C}_+[A, B]$ denote the set of continuous functions from $[A, B]$ to $\mathbb{R}^+$. Also, for $\alpha \leq (q+1)$, let $\mathcal{C}_+^\alpha[A, B] \subseteq \mathcal{C}_+[A, B]$ denote the set of functions that are α_0 times continuously differentiable, and for all $v_0 \in \mathcal{C}_+^\alpha[A, B]$, $\|v_0\|_\alpha < \infty$, where α_0 is largest integer less than or equals to α and the seminorm is defined by $\|v_0\|_\alpha = \sup_{X,X' \in [A,B], X \neq X'} \{|v_0^{(\alpha_0)}(X) - v_0^{(\alpha_0)}(X')|/|X - X'|^{\alpha - \alpha_0}\}$. The local support properties of B-splines make the models for the variance functions very flexible as is indicated by the following lemma.

Lemma 2. $\mathcal{C}_+^\alpha[A, B] \subseteq \mathcal{C}_+[A, B] \subseteq SN(\Pi_{\mathbf{V}})$.

Although technically the sup norm distance between linear combinations of B-splines and any continuous function can be made arbitrarily small by increasing the number of knots, for obvious reasons the actual bounds for the sup norm distance may not be very sharp if the function to be modeled is wiggly. However, for most applications of practical importance, the true variance function may be assumed to be smooth, that is, to belong to some $\mathcal{C}_+^\alpha[A, B]$ with $\alpha \geq 1$. Therefore, for practical reasons, it is only important that the smaller Hölder class of functions $\mathcal{C}_+^\alpha[A, B]$ belongs to the sup norm support of $\Pi_{\mathbf{V}}$. As shown in Section S.4.2 of the Supplementary Materials, the bounds for sup norm distance in this case will also be much sharper.

Since the models for the variance functions v_ℓ and the models for the density of the scaled errors $f_{\boldsymbol{\epsilon}}$ are separately very flexible, under model (16) on the measurement errors, the implied conditional and joint densities are also expected to be very flexible. This is investigated in the next lemma. For a given $\mathbf{X}$, let $\Pi_{\mathbf{U}|\mathbf{X}}$ denote the prior for $f_{\mathbf{U}|\mathbf{X}}$ induced by $\Pi_{\boldsymbol{\epsilon}}$ and $\Pi_{\mathbf{V}}$ under model (16). Define $\widetilde{\mathcal{F}}_{\mathbf{U}|\mathbf{X}} = \{f_{0\mathbf{U}|\mathbf{X}} : f_{0\mathbf{U}|\mathbf{X}}(\mathbf{U}) = \prod_{k=1}^p s_{0k}^{-1}(X_k) f_{0\boldsymbol{\epsilon}}\{\mathbf{S}_0^{-1}(\mathbf{X})\mathbf{U}\}, s_{0k}^2 \in \mathcal{C}_+[A_k, B_k] \text{ for } k = 1, \ldots, p, f_{0\boldsymbol{\epsilon}} \in \widetilde{\mathcal{F}}_{\boldsymbol{\epsilon}}\}$. Also let $\Pi_{\mathbf{U}|\mathbf{V}}$ denote the prior for the unknown conditional density of $\mathbf{U}$ induced by $\Pi_{\boldsymbol{\epsilon}}$ and $\Pi_{\mathbf{V}}$ under model (16). Define $\widetilde{\mathcal{F}}_{\mathbf{U}|\bullet} = \{f_{0\mathbf{U}|\bullet} : \text{ for any given } \mathbf{X} \in \mathcal{X},\ f_{0\mathbf{U}|\bullet} = f_{0\mathbf{U}|\mathbf{X}} \in \widetilde{\mathcal{F}}_{\mathbf{U}|\mathbf{X}}\}$. Finally, let $\Pi_{\mathbf{X},\mathbf{U}}$ denote the prior for the joint density of $(\mathbf{X}, \mathbf{U})$ induced by $\Pi_{\mathbf{X}}$, $\Pi_{\boldsymbol{\epsilon}}$ and $\Pi_{\mathbf{V}}$ under model (16). Define $\widetilde{\mathcal{F}}_{\mathbf{X},\mathbf{U}} = \{f_{0,\mathbf{X},\mathbf{U}} : f_{0,\mathbf{X},\mathbf{U}}(\mathbf{X}, \mathbf{U}) = f_{0,\mathbf{X}}(\mathbf{X}) f_{0,\mathbf{U}|\mathbf{X}}(\mathbf{U} \mid \mathbf{X}),$ where $f_{0\mathbf{X}} \in \widetilde{\mathcal{F}}_{\mathbf{X}}$ and $f_{0\mathbf{U}|\mathbf{X}} \in \widetilde{\mathcal{F}}_{\mathbf{U}|\mathbf{X}}$ for all $\mathbf{X} \in \mathcal{X}\}$.

Lemma 3. *1.* $\widetilde{\mathcal{F}}_{\mathbf{U}|\mathbf{X}} \subseteq KL(\Pi_{\mathbf{U}|\mathbf{X}})$ *for any given* $\mathbf{X} \in \mathcal{X}$.
2. For any $f_{0\mathbf{U}|\bullet} \in \widetilde{\mathcal{F}}_{\mathbf{U}|\mathbf{V}}$, $\Pi_{\mathbf{U}|\mathbf{V}}\{\sup_{\mathbf{X}\in\mathcal{X}} d_{KL}(f_{0\mathbf{U}|\mathbf{X}}, f_{\mathbf{U}|\mathbf{X}}) < \delta\} > 0$ *for all* $\delta > 0$.
3. $\widetilde{\mathcal{F}}_{\mathbf{X},\mathbf{U}} \subseteq KL(\Pi_{\mathbf{X},\mathbf{U}})$.

The flexibility of the implied model for the marginal density $f_{\mathbf{W}}$ is the subject of our final result. Since the only observed quantities are $\mathbf{W}_{ij}$, the support of the induced prior on $f_{\mathbf{W}}$ tells us about the types of likelihood functions the model can approximate.

Let $\Pi_{\mathbf{W}}$ denote the prior for the density of $\mathbf{W}$ induced by $\Pi_{\mathbf{X}}$, $\Pi_{\boldsymbol{\epsilon}}$ and $\Pi_{\mathbf{V}}$ under model (16). Also let $\widetilde{\mathcal{F}}_{\mathbf{W}} = \{f_{0\mathbf{W}} : f_{0\mathbf{W}}(\mathbf{W}) = \int f_{0\mathbf{X}}(\mathbf{X}) f_{0\mathbf{U}|\mathbf{X}}(\mathbf{W}-\mathbf{X}) d\mathbf{X}, f_{0\mathbf{X}} \in \widetilde{\mathcal{F}}_{\mathbf{X}}, f_{0\mathbf{U}|\bullet} \in \widetilde{\mathcal{F}}_{\mathbf{U}|\bullet}\}$, the class of densities $f_{0\mathbf{W}}$ that can be obtained as the convolution of two densities $f_{0\mathbf{X}}$ and $f_{0\mathbf{U}|\bullet}$, where $f_{0\mathbf{X}} \in \widetilde{\mathcal{F}}_{\mathbf{X}}$ and $f_{0\mathbf{U}|\bullet} \in \widetilde{\mathcal{F}}_{\mathbf{U}|\bullet}$.

Since the supports of $\Pi_{\mathbf{X}}$ and $\Pi_{\mathbf{U}|\mathbf{X}}$ are large, it is expected that the support of $\Pi_{\mathbf{W}}$ will also be large. However, because convolution is involved, investigation of KL support of $\Pi_{\mathbf{W}}$ is a difficult problem. A weaker but relevant concept is that of L_1 support. The L_1 distance between two densities f_0 and f, denoted by $||f_0 - f||_1$, is defined as $||f_0 - f||_1 = \int |f_0(Z) - f(Z)| dZ$. A density f_0 is said to belong to the L_1 support of Π_f if $\Pi_f(f : ||f_0 - f||_1 < \delta) > 0 \ \forall \delta > 0$. The class of densities in the L_1 support of Π_f is denoted by $L_1(\Pi_f)$. The following theorem shows that the L_1 support of $\Pi_{\mathbf{W}}$ is large.

Theorem 1. $\widetilde{\mathcal{F}}_{\mathbf{W}} \subseteq L_1(\Pi_{\mathbf{W}})$.

For space constraints, the technical details of the proofs of these results are deferred to Section S.4 of the Supplementary Materials. The proofs require that the number of mixture components K be allowed to vary over $\mathbb{N}$, the set of all positive integers, through priors, denoted by the generic $P_0(K)$, that assign positive probability to all $K \in \mathbb{N}$. Posterior computation of such methods will be computationally intensive, specially in a complicated multivariate set up like ours. To reduce computational complexity, in our implementation we keep the number of mixture components fixed at finite values. This does not necessarily come at the expense of the flexibility shown by the theoretical results of this section, since the priors $P_0(K)$ play only minor roles in the proofs. Indeed, the most important steps in the proofs of these results are actually to show that mixture models with fixed but sufficiently large number of components can approximate any target function with any desired level of accuracy provided they satisfy the above mentioned regularity conditions. With additional clauses to specify the desired degree of accuracy and lower bounds on the number of mixture components, the statements of the theoretical results presented in this section can be modified to eliminate the dependence on the priors $P_0(K)$. We have not done this to keep the statements of the theoretical results short and clean and, more importantly, to be able to clearly explain the trivial roles the priors $P_0(K)$ play in the proofs that strengthen our arguments in favor of using mixtures with finite fixed number of components. We have deferred these discussions to Section S.1 of the Supplementary Materials.

4 Choice of Hyper-Parameters

We discuss the choice of hyper-parameters in this section. To avoid unnecessary repetition, in this section and onwards, symbols sans the subscripts $\mathbf{X}$ and $\boldsymbol{\epsilon}$ are sometimes used as generics for similar components and parameters of the models. For example, K is a generic for $K_{\mathbf{X}}$ and $K_{\boldsymbol{\epsilon}}$; $\boldsymbol{\mu}_k$ is a generic for $\boldsymbol{\mu}_{\mathbf{X},k}$ and $\boldsymbol{\mu}_{\boldsymbol{\epsilon},k}$; and so on.

1. **Number of mixture components:** Practical application of our method requires that a decision be made on the number of mixture components $K_{\mathbf{X}}$ and $K_{\boldsymbol{\epsilon}}$ in the models for the densities $f_{\mathbf{X}}$ and $f_{\boldsymbol{\epsilon}}$, respectively.

Our simulation experiments suggest that when the true densities are finite mixtures of multivariate normals and $K_{\mathbf{X}}$ and $K_{\boldsymbol{\epsilon}}$ are assigned values greater than the corresponding true numbers, the MCMC chain often quickly reaches a steady state where the redundant components become empty. See Figures S.3, S.9 and S.10 in the Supplementary Materials for illustrations. These observations are similar to that made in the context of ordinary density estimation by Rousseau and Mengersen (2011) who studied the asymptotic behavior of the posterior for overfitted mixture models and showed that when $\alpha/K < L/2$, where L denotes the number of parameters specifying the component kernels, the posterior is stable and concentrates in regions with empty redundant components. We set $\alpha_{\mathbf{X}} = \alpha_{\boldsymbol{\epsilon}} = 1$ so that the condition $\alpha/K < L/2$ is satisfied.

Educated guesses about $K_{\mathbf{X}}$ and $K_{\boldsymbol{\epsilon}}$ may nevertheless be useful in safeguarding against gross overfitting that would result in a wastage of computation time and resources. The following simple strategies may be employed. Model based cluster analysis techniques as implemented by the mclust package in R (Fraley and Raftery, 2007) may be applied to the starting values of $\mathbf{X}_i$ and the corresponding residuals, obtained by fitting univariate submodels for each component of $\mathbf{X}$, to get some idea about $K_{\mathbf{X}}$ and $K_{\boldsymbol{\epsilon}}$. The chain may be started with larger values of $K_{\mathbf{X}}$ and $K_{\boldsymbol{\epsilon}}$ and after a few hundred iterations the redundant empty components may be deleted on the fly.

As shown in Section 3, our methods can approximate a much larger class of data generating densities, not just finite mixtures of multivariate normals, and even in those cases we found the strategy described above to be very effective. The parameter α now plays the role of a smoothing parameter, smaller values favoring a smaller number of mixture components and thus smoother densities. In simulation experiments involving multivariate t and multivariate Laplace distributions reported in the Supplemenatry Materials, and in some other cases not reported here, the values $\alpha_{\mathbf{X}} = \alpha_{\boldsymbol{\epsilon}} = 1$ worked well.

As we discuss in Section 7, the MIW method becomes highly numerically unstable when

the measurement errors are conditionally heteroscedastic and the true covariance matrices are highly sparse. In these cases in particular, the MIW method usually requires much larger sample sizes for the asymptotic results to hold and in finite samples the above mentioned strategy usually overestimates the required number of mixture components. See Figure S.2 in the Supplementary Materials for an illustration. Since mixtures based on $(K+1)$ components are at least as flexible as mixtures based on K components, as far as model flexibility is concerned, such overestimation is not an issue. But since this also results in clusters of smaller sizes, the estimates of the component specific covariance matrices become numerically even more unstable, further compounding the stability issues of the MIW model. In contrast, for the numerically more stable MLFA model, for the exact opposite reasons, the asymptotic results are valid for moderate sample sizes and such models are also more robust to overestimation of the number of nonempty clusters.

2. **Number of latent factors:** For the MLFA method, the MCMC algorithm summarized in Section 5 also requires that the component specific infinite factor models be truncated at some appropriate truncation level. The shrinkage prior again makes the model highly robust to overfitting allowing us to adopt a simple strategy. Since a latent factor characterization leads to a reduction in the number or parameters only when $q_k \leq \lceil (p+1)/2 \rceil$, where $\lceil s \rceil$ denotes the largest integer smaller than or equals to s, we simply set the truncation level at $q_k = q = \max\{2, \lceil (p+1)/2 \rceil\}$ for all the components. We also experimented by setting the truncation level at $q_k = q = p$ for all k with the results remaining practically the same. The shrinkage prior, being continuous in nature, does not set the redundant columns to exact zeroes, but it adaptively shrinks the redundant parameters sufficiently towards zero, thus producing stable and efficient estimates of the densities being modeled.

3. **Other hyper-parameters:** We take an empirical Bayes type approach to assign values to other hyper-parameters. We set $\boldsymbol{\mu}_{\mathbf{X},0} = \overline{\mathbf{X}}^{(0)}$, the overall mean of $\mathbf{X}^{(0)}_{1:n}$, where $\mathbf{X}^{(0)}_{1:n}$ denote the starting values of $\mathbf{X}_{1:n}$ for the MCMC sampler discussed in Section 5. For the scaled errors we set $\boldsymbol{\mu}_{\boldsymbol{\epsilon},0} = \mathbf{0}$. For the MIW model we take $\nu_0 = (p+2)$, the smallest possible integral value of ν_0 for which the prior mean of $\boldsymbol{\Sigma}_k$ exists. We then take $\boldsymbol{\Sigma}_{\mathbf{X},0}/2 = \boldsymbol{\Psi}_{\mathbf{X},0} = \text{cov}(\overline{\mathbf{X}}^{(0)}_{1:n})$. These choices imply $E(\boldsymbol{\Sigma}_{\mathbf{X},k}) = \boldsymbol{\Psi}_{\mathbf{X},0} = \text{cov}(\overline{\mathbf{X}}^{(0)})$ and, since the variability of each component is expected to be significantly less than the overall variability, ensure noninformativeness. Similarly, for the scaled errors we take $\boldsymbol{\Sigma}_{\boldsymbol{\epsilon},0}/2 = \boldsymbol{\Psi}_{\boldsymbol{\epsilon},0} = \text{cov}(\boldsymbol{\epsilon}^{(0)}_{1:N})$. For the MLFA model, the hyper-parameters specifying the prior for $\boldsymbol{\Lambda}$ are set at $a_1 = 1, a_h = 2$ for all $h \geq 2$, and $\nu = 1$. Inverse gamma priors with parameters $a_\sigma = 1.1, b_\sigma = 1$ are placed on the elements of $\boldsymbol{\Omega}$. For each k, the variance functions were modeled using quadratic (q=2) B-splines based on $(2 \times 2 + 5 + 1) = 10$ equidistant knot points on $[A_k, B_k] = [\min(\overline{\mathbf{W}}_{k,1:n}) - 0.1 \text{ range}(\overline{\mathbf{W}}_{k,1:n}), \max(\overline{\mathbf{W}}_{k,1:n}) + 0.1 \text{ range}(\overline{\mathbf{W}}_{k,1:n})]$, where $\overline{\mathbf{W}}_{\ell,1:n}$ denotes the subject specific means corresponding to ℓ^{th} component.

5 Posterior Computation

Samples from the posterior can be drawn using Gibbs sampling techniques. In what follows $\boldsymbol{\zeta}$ denotes a generic variable that collects the observed proxies $\mathbf{W}_{1:N}$ and all the parameters of a model, including the imputed values of $\mathbf{X}_{1:n}$ and $\boldsymbol{\epsilon}_{1:N}$, that are not explicitly mentioned.

Carefully chosen starting values can facilitate convergence of the sampler. The posterior means of the $X_{i\ell}$'s, obtained by fitting univariate versions of the proposed multivariate methods for each component of $\mathbf{X}$, are used as the starting values for the multivariate sampler. The number of mixture components are initialized at $K_{\mathbf{X}} = (m_{\mathbf{X}} + 2)$, where $m_{\mathbf{X}}$ denotes the optimal number of clusters returned by model based clustering algorithm implemented by the mclust package in R applied to the corresponding initial values $\mathbf{X}^{(0)}_{1:n}$. The component specific mean vectors of the nonempty clusters are set at the mean of $\mathbf{X}^{(0)}_i$ values that belong to that cluster. The component specific mean vectors of the two empty clusters are set at $\overline{\mathbf{X}}^{(0)}$, the overall mean of $\mathbf{X}^{(0)}_{1:n}$. For the MIW model, the initial values of the cluster specific covariance matrices are chosen in a similar fashion. The mixture probabilities for the k^{th} nonempty cluster is set at $\boldsymbol{\pi}_{\mathbf{X},k} = n_k/n$, where n_k denotes the number of $\mathbf{X}^{(0)}_i$ belonging to the k^{th} cluster. The mixture probabilities of the empty clusters are initialized at zero. For the MLFA method, the starting values of all elements of $\boldsymbol{\Lambda}$ and $\boldsymbol{\eta}$ are set at zero. The starting values for the elements of $\boldsymbol{\Omega}$ are chosen to equal the variances of the corresponding starting values. The parameters specifying the density of the scaled errors are initialized in a similar manner. The MCMC iterations comprise the following steps.

1. **Updating the parameters specifying $f_{\mathbf{X}}$:** For the MIW model the parameters specifying the density $f_{\mathbf{X}}$ are updated using the following steps.

$$
\begin{aligned}
(\boldsymbol{\pi} \mid \boldsymbol{\zeta}) &\sim \text{Dir}(\alpha/K + n_1, \alpha/K + n_2, \ldots, \alpha/K + n_K) \\
(C_i \mid \boldsymbol{\zeta}) &\sim \text{Mult}(1, p_{i1}, p_{i2}, \ldots, p_{iK}), \\
(\boldsymbol{\mu}_k \mid \boldsymbol{\zeta}) &\sim \text{MVN}_p(\boldsymbol{\mu}^{(n)}_k, \boldsymbol{\Sigma}^{(n)}_k), \\
(\boldsymbol{\Sigma}_k \mid \boldsymbol{\zeta}) &\sim \text{IW}_p\{n_k + \nu_0, \textstyle\sum_{i:C_i=k}(\mathbf{X}_i - \boldsymbol{\mu}_k)(\mathbf{X}_i - \boldsymbol{\mu}_k)^{\mathrm{T}} + \boldsymbol{\Psi}_0\},
\end{aligned}
$$

where $n_k = \sum_i 1(C_i = k)$, $p_{ik} \propto \pi_k \times \text{MVN}_p(\mathbf{X}_i \mid \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$, $\boldsymbol{\Sigma}^{(n)}_k = (\boldsymbol{\Sigma}^{-1}_0 + n_k\boldsymbol{\Sigma}^{-1}_k)^{-1}$ and $\boldsymbol{\mu}^{(n)}_k = \boldsymbol{\Sigma}^{(n)}_k \left\{\boldsymbol{\Sigma}^{-1}_k \sum_{i:C_i=k} \mathbf{X}_i + \boldsymbol{\Sigma}^{-1}_0\boldsymbol{\mu}_0\right\}$. To update the parameters specifying the covariance matrices in the MLFA model, the sampler cycles through the following steps.

$$
\begin{aligned}
(\boldsymbol{\lambda}_{k,j} \mid \boldsymbol{\zeta}) &\sim \text{MVN}_q\{(\mathbf{D}^{-1}_{k,j} + \sigma_j^{-2}\boldsymbol{\eta}^{\mathrm{T}}_k\boldsymbol{\eta}_k)^{-1}\sigma_j^{-2}\boldsymbol{\eta}^{\mathrm{T}}_k(\mathbf{X}^{(j)}_k - \boldsymbol{\mu}^{(j)}_k), (\mathbf{D}^{-1}_{k,j} + \sigma_j^{-2}\boldsymbol{\eta}^{\mathrm{T}}_k\boldsymbol{\eta}_k)^{-1}\}, \\
(\boldsymbol{\eta}_i \mid C_i = k, \boldsymbol{\zeta}) &\sim \text{MVN}_q\{(\mathrm{I}_q + \boldsymbol{\Lambda}^{\mathrm{T}}_k\boldsymbol{\Omega}^{-1}\boldsymbol{\Lambda}_k)^{-1}\boldsymbol{\Lambda}^{\mathrm{T}}_k\boldsymbol{\Omega}^{-1}(\mathbf{X}_i - \boldsymbol{\mu}_k), (\mathrm{I}_q + \boldsymbol{\Lambda}^{\mathrm{T}}_k\boldsymbol{\Omega}^{-1}\boldsymbol{\Lambda}_k)^{-1}\}, \\
(\sigma_j^2 \mid \boldsymbol{\zeta}) &\sim \text{Inv-Ga}\left\{a_\sigma + n/2, b_\sigma + (1/2)\textstyle\sum_{i=1}^n (X_{ij} - \boldsymbol{\mu}_{C_i,j} - \boldsymbol{\lambda}^{\mathrm{T}}_{C_i,j}\boldsymbol{\eta}_i)^2\right\}, \\
(\phi_{k,jh} \mid \boldsymbol{\zeta}) &\sim \text{Ga}\{(\nu+1)/2, (\nu + \tau_{k,h}\lambda^2_{k,jh})/2\}, \\
(\delta_{k,h} \mid \boldsymbol{\zeta}) &\sim \text{Ga}\{a_h + p(q-h+1)/2, 1 + \textstyle\sum_{\ell=1}^q \tau^{(h)}_{k,\ell} \sum_{j=1}^p \phi_{k,j\ell}\lambda^2_{k,j\ell}/2\},
\end{aligned}
$$

where $D_{k,j}^{-1} = \text{diag}(\phi_{k,j1}\tau_{k,1}, \ldots, \phi_{k,jq}\tau_{k,q})$, $\tau_{k,\ell}^{(h)} = \prod_{t=1,t\neq h}^{\ell} \delta_{k,t}$, $\mathbf{X}_k^{(j)} = (X_{i_1 j}, X_{i_2 j}, \ldots, X_{i_{n_k} j})^{\mathrm{T}}$, $\boldsymbol{\eta}_k^{n_k \times q} = (\boldsymbol{\eta}_{i_1}, \boldsymbol{\eta}_{i_2}, \ldots, \boldsymbol{\eta}_{i_{n_k}})^{\mathrm{T}}$, $\{i_1, i_2, \ldots, i_{n_k}\} = \{i : C_i = k\}$.

2. **Updating the parameters specifying $f_{\boldsymbol{\epsilon}}$:** Except for one simple additional step that enforces the zero mean restriction, the steps to update the parameters specifying the error distribution are similar and thus excluded.

3. **Updating the values of X:** When the measurement errors are independent of $\mathbf{X}$, the $\mathbf{X}_i$ have closed form full conditionals given by

$$(\mathbf{X}_i \mid C_{\mathbf{X},i} = k, C_{\boldsymbol{\epsilon},i1} = k_1, \ldots, C_{\boldsymbol{\epsilon},im_i} = k_{m_i}, \boldsymbol{\zeta}) \;\sim\; \text{MVN}_p(\boldsymbol{\mu}_{\mathbf{X}}^{(n)}, \boldsymbol{\Sigma}_{\mathbf{X}}^{(n)}),$$

where $\boldsymbol{\Sigma}_{\mathbf{X}}^{(n)} = (\boldsymbol{\Sigma}_{\mathbf{X},k}^{-1} + \sum_{j=1}^{m_i} \boldsymbol{\Sigma}_{\boldsymbol{\epsilon},k_j}^{-1})^{-1}$ and $\boldsymbol{\mu}_{\mathbf{X}}^{(n)} = \boldsymbol{\Sigma}_{\mathbf{X}}^{(n)}(\boldsymbol{\Sigma}_{\mathbf{X},k}^{-1}\boldsymbol{\mu}_{\mathbf{X},k} + \sum_{j=1}^{m_i} \boldsymbol{\Sigma}_{\boldsymbol{\epsilon},k_j}^{-1}\mathbf{W}_{ij})$. For conditionally heteroscedastic measurement errors, the full conditionals are given by

$$\begin{aligned}&(\mathbf{X}_i \mid C_{\mathbf{X},i} = k, C_{\boldsymbol{\epsilon},i1} = k_1, \ldots, C_{\boldsymbol{\epsilon},im_i} = k_{m_i}, \boldsymbol{\zeta}) \\ &\quad \propto \text{MVN}_p(\mathbf{X}_i \mid \boldsymbol{\mu}_{\mathbf{X},k}, \boldsymbol{\Sigma}_{\mathbf{X},k}) \times \textstyle\prod_{j=1}^{m_i} \text{MVN}_p\{\mathbf{W}_{ij} \mid \mathbf{X}_i + \mathbf{S}(\mathbf{X}_i)\boldsymbol{\mu}_{\boldsymbol{\epsilon},k_j}, \mathbf{S}(\mathbf{X}_i)\boldsymbol{\Sigma}_{\boldsymbol{\epsilon},k_j}\mathbf{S}(\mathbf{X}_i)\},\end{aligned}$$

The full conditionals do not have closed forms. Metropolis-Hastings (MH) steps with multivariate truncated normal proposals are used within the Gibbs sampler.

4. **Updating the parameters specifying s_ℓ:** When the measurement errors are conditionally heteroscedastic, we first estimate the variance functions $s_\ell^2(X_{i\ell})$ by fitting univariate submodels $W_{ij\ell} = X_{i\ell} + s_\ell(X_{i\ell})\epsilon_{ij\ell}$ for each k. The details are provided in Section 6. The parameters characterizing other components of the full model are then sampled using the Gibbs sampler described above, keeping the estimates of the variance functions fixed.

6 Estimation of the Variance Functions

When the measurement errors are conditionally heteroscedastic, we need to update the parameters $\boldsymbol{\xi}_k$ that specify the variance functions $s_\ell^2(X_{i\ell})$. These parameters do not have closed form full conditionals. MCMC algorithms, where we tried to integrate MH steps for $\boldsymbol{\xi}_k$ with the sampler for the parameters specifying $f_{\boldsymbol{\epsilon}}$, were numerically unstable and failed to converge sufficiently quickly. We need to supply the values of the scaled errors to step 2 of the algorithm described in Section 5 and the instability stems from the operation $\boldsymbol{\epsilon}_{ij} = \mathbf{S}(\mathbf{X}_i)^{-1}\mathbf{U}_{ij}$ required to calculate these values.

To solve the problem, we adopt a novel two-stage procedure. First, for each k, we estimate the functions $s_\ell^2(X_{i\ell})$ by fitting the univariate submodels $W_{ij\ell} = X_{i\ell} + s_\ell(X_{i\ell})\epsilon_{ij\ell}$. The problem of numerical instability arising out of the operation to determine the values of the scaled errors remains in these univariate subproblems too. But the following lemma from Pelenis (2014), presented here for easy reference, provides us with an escape route by allowing us to avoid this operation in the first place.

Lemma 4. *Let* $\boldsymbol{\theta}_{1:K} = \{(\pi_k, \mu_k, \sigma_k^2)\}_{k=1}^K$ *be such that*

$$f_1(\epsilon \mid \boldsymbol{\theta}_{1:K}) = \sum_{k=1}^K \pi_k \; Normal(\epsilon \mid \mu_k, \sigma_k^2), \quad \text{with} \quad \sum_{k=1}^K \pi_k = 1, \quad \sum_{k=1}^K \pi_k \mu_k = 0. \tag{19}$$

Then there exists a set of parameters $\boldsymbol{\theta}^\star_{1:(K-1)}$ *such that*

$$\begin{aligned} f_1(\epsilon \mid \boldsymbol{\theta}_{1:K}) = f_2(\epsilon \mid \boldsymbol{\theta}^\star_{1:(K-1)}) = \textstyle\sum_{k=1}^{K-1} \pi_k^\star \; \sum_{r=1}^2 p^\star_{k,r} Normal(\epsilon \mid \mu^\star_{k,r}, \sigma^{\star 2}_{k,r}), \\ \textstyle\sum_{k=1}^{K-1} \pi_k^\star = 1, \quad \sum_{r=1}^2 p^\star_{k,r} = 1, \quad \sum_{r=1}^2 p^\star_{k,r} \mu^\star_{k,r} = 0 \; \forall k. \end{aligned} \tag{20}$$

Lemma 4 implies that the univariate submodels for the density of the scaled errors given by (19) has a reparametrization (20) where each component is itself a two-component normal mixture with its mean restricted at zero. The reparametrization (20) thus replaces the zero mean restriction on (19) by similar restrictions on each of its components. These restrictions also imply that each mixture component in (20) can be further reparametrized by only four free parameters. One such parametrization could be in terms of $\widetilde{\boldsymbol{\theta}}_k = (\widetilde{p}_k, \widetilde{\mu}_k, \widetilde{\sigma}^2_{k,1}, \widetilde{\sigma}^2_{k,2})$, where $(p^\star_{k,1}, \sigma^{\star 2}_{k,1}, \sigma^{\star 2}_{k,2}) = (\widetilde{p}_k, \widetilde{\sigma}^2_{k,1}, \widetilde{\sigma}^2_{k,2})$ and $\mu^\star_{k,r} = c_{k,r}\widetilde{\mu}_k$, where $c_{k,1} = (1-\widetilde{p}_k)/\{\widetilde{p}_k^2 + (1-\widetilde{p}_k)^2\}^{1/2}$ and $c_{k,2} = -\widetilde{p}_k/\{\widetilde{p}_k^2 + (1-\widetilde{p}_k)^2\}^{1/2}$. Letting p_0 denote the prior assigned to $\widetilde{\boldsymbol{\theta}}_k$, the full conditional of $\widetilde{\boldsymbol{\theta}}_k$ in terms of the conditional likelihood $f_{U|X}$ is proportional to $p_0(\widetilde{\boldsymbol{\theta}}_k) \prod_{ij:C_{\epsilon,ij\ell}=k} f_{U|X}(U_{ij\ell} \mid X_{i\ell}, \boldsymbol{\xi}_\ell, \widetilde{\boldsymbol{\theta}}_k, \boldsymbol{\zeta})$. The problem of numerical instability can now be tackled by using MH steps to update not only the parameters $\boldsymbol{\xi}_\ell$ specifying the variance functions but also the parameters $\{\widetilde{\boldsymbol{\theta}}_k\}_k$ characterizing the density f_ϵ using the conditional likelihood $f_{U|X}$ (and not f_ϵ itself), thus escaping the need to separately determine the values of the scaled errors.

The model for the density of the measurement errors in Sarkar, et al. (2014) was essentially an infinite dimensional extension of this reparametrized finite dimensional submodel. The priors and the hyper-parameters for the univariate submodels may be chosen following their suggestions. The strategy of exploiting the properties of overfitted mixture models to determine the number of mixture components described in Section 4 can also be applied to the univariate subproblems. High precision estimates of the variance functions can be obtained using these reparametrized finite dimensional univariate deconvolution models. See Figure 1 and also Figures S.4 and S.13 in the Supplementary Materials for illustrations.

A similar reparametrization exists for the multivariate problem too, but the strategy would not be very effective in a multivariate set up as it would require updating the mean vectors and the covariance matrices involved in $f_{\boldsymbol{\epsilon}}$ through MH steps which are not efficient in simultaneous updating of large numbers of parameters. After estimating the functions variance functions from the univariate submodels, we therefore keep these estimates fixed and sample other parameters using the Gibbs sampler described in Section 5.

7 Simulation Experiments

The mean integrated squared error (MISE) of estimation of $f_{\mathbf{X}}$ by $\widehat{f}_{\mathbf{X}}$ is defined as $MISE = E\int\{f_{\mathbf{X}}(\mathbf{X}) - \widehat{f}_{\mathbf{X}}(X)\}^2 dX$. Based on B simulated data sets, a Monte Carlo estimate of MISE is given by $MISE_{est} = B^{-1}\sum_{b=1}^{B}\sum_{m=1}^{M}\{f_{\mathbf{X}}(\mathbf{X}_{b,m}) - \widehat{f}_{\mathbf{X}}^{(b)}(\mathbf{X}_{b,m})\}^2/p_0(\mathbf{X}_{b,m})$, where $\{\mathbf{X}_{b,m}\}_{b=1,m=1}^{B,M}$ are random samples from the density p_0. We designed simulation experiments to evaluate the MISE performance of the proposed models for a wide range of possibilities. The MISEs we report here are all based on 100 simulated data sets and $M = 10^6$ samples generated from each of the two densities (a) $p_0 = f_{\mathbf{X}}$, the true density of $\mathbf{X}$, and (b) p_0 that is uniform on the hypercube with edges $\min_k\{\boldsymbol{\mu}_{\mathbf{X},k} - 3\mathbf{1}_p\}$ and $\max_k\{\boldsymbol{\mu}_{\mathbf{X},k} + 3\mathbf{1}_p\}$. With carefully chosen initial values and proposal densities for the MH steps, we were able to achieve quick convergence for the MCMC samplers. The use of exchangeable Dirichlet priors helped simplify mixing issues (Geweke, 2007). See Section S.1.1 in the Supplementary Materials for additional discussions. We programmed our methods in R. In each case, we ran 3000 MCMC iterations and discarded the initial 1000 iterations as burn-in. The post burn-in samples were thinned by a thinning interval of length 5. For the univariate samplers, 1000 MCMC iterations with a burn-in of 500 sufficed to produce stable estimates of the variance functions. In our experiments with much larger iteration numbers and burn-ins, the MISE performances remained practically the same. This being the first article that tries to solve the problem of multivariate density deconvolution when the measurement error density is unknown, the proposed MIW and MLFA models have no competitors. We thus compared our models with a naive Bayesian method that ignores measurement errors and treats the subject specific means as precisely measured observations instead, modeling $f_{\mathbf{X}}$ by a finite mixture of multivariate normals as in (2) with inverse Wishart priors on the component specific covariance matrices.

We considered two choices for the sample size $n = 500, 1000$. For each subject, we simulated $m_i = 3$ replicates. The true density of $\mathbf{X}$ was chosen to be $f_{\mathbf{X}}(\mathbf{X}) = \sum_{k=1}^{K_{\mathbf{X}}}\pi_{\mathbf{X},k}\ \mathrm{MVN}_p(\mathbf{X} \mid \boldsymbol{\mu}_{\mathbf{X},k}, \boldsymbol{\Sigma}_{\mathbf{X},k})$ with $p = 4$, $K_{\mathbf{X}} = 3$, $\boldsymbol{\pi}_{\mathbf{X}} = (0.25, 0.50, 0.25)^{\mathrm{T}}$, $\boldsymbol{\mu}_{\mathbf{X},1} = (0.8, 6, 4, 5)^{\mathrm{T}}$, $\boldsymbol{\mu}_{\mathbf{X},2} = (2.5, 4, 5, 6)^{\mathrm{T}}$ and $\boldsymbol{\mu}_{\mathbf{X},3} = (6, 4, 2, 4)^{\mathrm{T}}$. For the density of the measurement errors $f_{\boldsymbol{\epsilon}}$ we considered two choices, namely

1. $f_{\boldsymbol{\epsilon}}^{(1)}(\boldsymbol{\epsilon}) = \mathrm{MVN}_p(\boldsymbol{\epsilon} \mid \mathbf{0}, \boldsymbol{\Sigma}_{\boldsymbol{\epsilon}})$, and
2. $f_{\boldsymbol{\epsilon}}^{(2)}(\boldsymbol{\epsilon}) = \sum_{k=1}^{K_{\boldsymbol{\epsilon}}}\pi_{\boldsymbol{\epsilon},k}\ \mathrm{MVN}_p(\boldsymbol{\epsilon} \mid \boldsymbol{\mu}_{\boldsymbol{\epsilon},k}, \boldsymbol{\Sigma}_{\boldsymbol{\epsilon},k})$ with $K_{\boldsymbol{\epsilon}} = 3$, $\boldsymbol{\pi}_{\boldsymbol{\epsilon}} = (0.2, 0.6, 0.2)^{\mathrm{T}}$, $\boldsymbol{\mu}_{\boldsymbol{\epsilon},1} = (-0.3, 0, 0.3, 0)^{\mathrm{T}}$, $\boldsymbol{\mu}_{\boldsymbol{\epsilon},2} = (-0.5, 0.4, 0.5, 0)^{\mathrm{T}}$ and $\boldsymbol{\mu}_{\boldsymbol{\epsilon},3} = -(\pi_{\boldsymbol{\epsilon},1}\boldsymbol{\mu}_{\boldsymbol{\epsilon},1} + \pi_{\boldsymbol{\epsilon},2}\boldsymbol{\mu}_{\boldsymbol{\epsilon},2})/\pi_{\boldsymbol{\epsilon},3}$.

For the component specific covariance matrices, we set $\boldsymbol{\Sigma}_{\mathbf{X},k} = \mathbf{D}_{\mathbf{X}}\boldsymbol{\Sigma}_{\mathbf{X},0}\mathbf{D}_{\mathbf{X}}$ for each k, where $\mathbf{D}_{\mathbf{X}} = \mathrm{diag}(0.75^{1/2}, \ldots, 0.75^{1/2})$. Similarly, $\boldsymbol{\Sigma}_{\boldsymbol{\epsilon},k} = \mathbf{D}_{\boldsymbol{\epsilon}}\boldsymbol{\Sigma}_{\boldsymbol{\epsilon},0}\mathbf{D}_{\boldsymbol{\epsilon}}$ for each k, where $\mathbf{D}_{\boldsymbol{\epsilon}} = \mathrm{diag}(0.3^{1/2}, \ldots, 0.3^{1/2})$. For each pair of $f_{\mathbf{X}}$ and $f_{\boldsymbol{\epsilon}}$, we considered four types of covariance structures for $\boldsymbol{\Sigma}_{\mathbf{X},0} = \{(\sigma_{ij}^{\mathbf{X},0})\}$ and $\boldsymbol{\Sigma}_{\boldsymbol{\epsilon},0} = \{(\sigma_{ij}^{\boldsymbol{\epsilon},0})\}$, namely

1. Identity (I): $\boldsymbol{\Sigma}_{\mathbf{X},0} = \boldsymbol{\Sigma}_{\boldsymbol{\epsilon},0} = \mathrm{I}_p$,

2. Latent Factor (LF): $\boldsymbol{\Sigma}_{\mathbf{X},0} = \boldsymbol{\Lambda}_{\mathbf{X}}\boldsymbol{\Lambda}_{\mathbf{X}} + \boldsymbol{\Omega}_{\mathbf{X}}$, with $\boldsymbol{\Lambda}_{\mathbf{X}} = (0.7, \ldots, 0.7)^{\mathrm{T}}$ and $\boldsymbol{\Omega}_{\mathbf{X}} = \mathrm{diag}(0.51, \ldots, 0.51)$, and $\boldsymbol{\Sigma}_{\boldsymbol{\epsilon},0} = \boldsymbol{\Lambda}_{\boldsymbol{\epsilon}}\boldsymbol{\Lambda}_{\boldsymbol{\epsilon}} + \boldsymbol{\Omega}_{\boldsymbol{\epsilon}}$, with $\boldsymbol{\Lambda}_{\boldsymbol{\epsilon}} = (0.5, \ldots, 0.5)^{\mathrm{T}}$ and $\boldsymbol{\Omega}_{\boldsymbol{\epsilon}} = \mathrm{diag}(0.75, \ldots, 0.75)$,

3. Autoregressive (AR): $\sigma_{ij}^{\mathbf{X},0} = 0.7^{|i-j|}$ and $\sigma_{ij}^{\boldsymbol{\epsilon},0} = 0.5^{|i-j|}$ for each (i, j), and

4. Exponential (EXP): $\sigma_{ij}^{\mathbf{X},0} = \exp(-0.5\,|i-j|)$ and $\sigma_{ij}^{\boldsymbol{\epsilon},0} = \exp(-0.9\,|i-j|)$ for each (i, j).

The parameters were chosen to produce a wide variety of one and two dimensional marginal densities, see Figure 3 and also Figure 5 in the Supplementary Materials. Scale adjustments by multiplication with $\mathbf{D}_{\mathbf{X}}$ and $\mathbf{D}_{\boldsymbol{\epsilon}}$ were done so that the simulated values of each component of $\mathbf{X}$ fall essentially in the range $(-2, 6)$ and the simulated values of all components of $\boldsymbol{\epsilon}$ fall essentially in the range $(-3, 3)$. For conditionally heteroscedastic measurement errors, we set the true variance functions at $s_\ell^2(X) = (1 + X/4)^2$ for each component k. A total of 16 ($2 \times 1 \times 2 \times 4$) cases were thus considered for both independent and conditionally heteroscedastic measurement errors.

We first discuss the results of the simulation experiments when the measurement errors $\mathbf{U}$ were independent of $\mathbf{X}$. The estimated MISEs are presented in Table 1. When the true $f_{\boldsymbol{\epsilon}}$ was a single component multivariate normal, the MLFA model produced the lowest MISE when the true covariance matrices were diagonal. In all other cases the MIW model produced the best results. When the true $f_{\boldsymbol{\epsilon}}$ was a mixture of multivariate normals, the model complexity increases and the performance of the MIW model started to deteriorate. In this case, the MLFA model dominated the MIW model when the true covariance matrices were either diagonal or had a latent factor characterization.

The estimated MISEs for the cases when $\mathbf{U}$ were conditionally heteroscedastic are presented in Table 2. Models that accommodate conditional heteroscedasticity are significantly more complex compared to models that assume independence of the measurement errors from $\mathbf{X}$. The numerically more stable MLFA model thus out-performed the MIW model in all 32 cases. The improvements were particularly significant when the true covariance matrices were sparse and the number of subjects was small ($n = 500$). The true and estimated univariate and bivariate marginals of $f_{\mathbf{X}}$ produced by the MIW and the MLFA methods when the true density of the scaled errors was a mixture of multivariate normals ($f_{\boldsymbol{\epsilon}}^{(2)}$) and the component specific covariance matrices were diagonal (I) are summarized in Figure 2 and Figure 3, respectively. The true and estimated univariate and bivariate marginals for the density of the scaled errors $f_{\boldsymbol{\epsilon}}$ for this case produced by the two methods are summarized in Figure 4 and Figure 5, respectively. The true and the estimated variance functions produced by the univariate submodels are summarized in Figure 1. Comparisons between

Figure 2 and Figure 3 illustrate the limitations of the MIW models in capturing high dimensional sparse covariance matrices and the improvements that can be achieved by the MLFA models. The estimates of $f_{\boldsymbol{\epsilon}}$ produced by the two methods are in better agreement. This may be attributed to the fact that many more residuals are available for estimating $f_{\boldsymbol{\epsilon}}$ than there are $\mathbf{X}_i$'s to estimate $f_{\mathbf{X}}$. Figure 1 in the main paper and Figures S.4 and S.13 in the Supplementary Materials show that the univariate submodels can recover the true variance functions well. Additional figures when the true covariance matrices had auto-regressive structure (AR) are presented in the Supplementary Materials. In this case the true covariance matrices were not sparse. The MLFA method still vastly dominated the MIW method when the sample size was small ($n = 500$). When the sample size was large ($n = 1000$) the two methods produced comparable results.

The proposed deconvolution methods, in particular the MLFA method, are highly scalable. In small scale simulations, not reported here, we tried $p = 6, 8$ and 10 and observed good empirical performance. We have focused here on $p = 4$ dimensional problems since with $p = 4$ the numbers of univariate and bivariate marginals, $p = 4$ and $\binom{p}{2} = 6$, remain manageable and the results are conveniently graphically summarized.

As shown by the results of Section 3, the proposed deconvolution models can approximate a large class of data generating densities, not just finite mixtures of multivariate normals. To illustrate this, additional simulations for multivariate t and multivariate Laplace distributions with similar MISE patterns are presented in the Supplementary Materials.

8 Example

Dietary habits are known to be leading causes of many chronic diseases. Accurate estimation of the distributions of dietary intakes is thus important in nutritional epidemiologic surveillance and epidemiology. Nutritionists are typically interested not just in the consumption patterns of individual dietary components but also in their joint consumption patterns. By the very nature of the problem, $\mathbf{X}$, the average long term daily intakes of the dietary components, can never be directly observed. Data are thus typically collected from a representative sample of the population in the form of dietary recalls, the subjects participating in the study remembering and reporting the type and amount of food they had consumed in the past 24 hours. The problem of estimating the joint consumption pattern of the dietary components from the contaminated 24-hour recalls then becomes a problem of multivariate density deconvolution.

A large scale epidemiologic study conducted by the National Cancer Institute, the Eating at America's Table (EATS) study (Subar, et al., 2001), serves as the motivation for this paper. In this study $n = 965$ participants were interviewed $m_i = 4$ times over the course of

a year and their 24 hour dietary recalls ($\mathbf{W}_{ij}$'s) were recorded. The goal is to estimate the joint consumption patterns of the true daily intakes ($\mathbf{X}_i$'s).

To illustrate our methodology, we consider the problem of estimating the joint consumption pattern of four dietary components, namely carbohydrate (a), fiber (b), protein (c) and a mineral potassium (d). Figure 6 shows the plots of subject-specific means versus subject-specific variances for daily intakes of the dietary components with the estimates of the variance functions produced by univariate submodels superimposed over them. As is clearly identifiable from this plot, conditional heteroscedasticity is a very prominent feature of the measurements errors contaminating the 24 hour recalls. The estimated univariate and bivariate marginal densities of average long term daily intakes of the dietary components produced by the MIW method and the MLFA method are summarized in Figure 7. The estimated univariate and bivariate marginal densities for the scaled errors are summarized in Figure 8. The estimated marginals of $\mathbf{X}$ produced by the two methods look quite different, while the estimated marginals of $\boldsymbol{\epsilon}$ are in close agreement. The estimated univariate and bivariate marginal densities of the long term intakes of the dietary components produced by the MIW model look irregular and unstable, whereas the estimates produced by the MLFA model look relatively more regular and stable. In experiments not reported here, we observed that the estimates produced by the MIW method were sensitive to the choice of the number of mixture components, but the estimates produced by the MLFA model were quite robust. The trace plots and the frequency distributions of the of the numbers of nonempty mixture components are summarized in Figures S.11 and S.12 in the Supplementary Materials and provide some idea about the relative stability of the two methods. These observations are similar to that made in Section 7 for conditionally heteroscedastic measurement errors and sparse covariance matrices.

We next comment only on the estimates produced by the MLFA method assuming them to be closer to the truth. The estimates show that the long term daily intakes of the four dietary components are strongly correlated. The shapes of the bivariate consumption patterns suggest deviations from normality. Similarly, the shapes of the bivariate marginals for the scaled errors suggest that the measurement errors in the reported 24 hour recalls are positively correlated and deviate from normality. People who consume more are expected to do so for most dietary components. Strong correlations between the intakes of the dietary components are thus somewhat expected. The correlations among different components of the measurement errors suggest that people usually have a tendency to either over-report or under-report the daily intakes. These findings illustrate the importance of robust but numerically stable multivariate deconvolution methods in nutritional epidemiologic studies.

9 Discussion

We considered the problem of multivariate density deconvolution when the measurement error density is not known but replicated proxies are available for some individuals. We used flexible finite mixtures of multivariate normal kernels with symmetric Dirichlet priors on the mixture probabilities to model both the density of interest and the density of the measurement errors. We proposed a novel technique to make the model for the density of the errors satisfy a zero mean restriction. We showed that the dense parametrization of inverse Wishart priors are not suitable for modeling covariance matrices in the presence of measurement errors. We proposed a numerically more stable approach based on latent factor characterization of the covariance matrices with sparsity inducing priors on the factor loading matrices. We built models for conditionally heteroscedastic measurement errors taking into account their somewhat unique characteristics and proposed a novel two stage procedure to tackle the computational obstacles. Being the first article that allows the density of the measurement errors to be unknown and free from parametric laws and also accommodates conditional heteroscedasticity, this article, we believe, makes an important contribution to the sparse literature on the problem of multivariate density deconvolution.

The methodological contributions of this article are not limited to deconvolution problems. The mixture of latent factor analyzers with sparsity inducing priors on the factor loading matrices can be used in other high dimensional applications including ordinary density estimation. The techniques proposed in Section 2.2.1 to enforce the mean zero moment restriction on the measurement errors can be readily used to model multivariate regression errors that are distributed independently of the predictors. As commented upon in Section 2.2.2 and detailed in the Supplementary Materials, the covariance regression techniques developed in Hoff and Niu (2012) and Fox and Dunson (2013) for modeling conditionally varying regression errors are not suitable for modeling conditionally heteroscedastic measurement errors. Similarly, the models for conditionally heteroscedastic multivariate measurement errors developed in this article are not suitable for modeling conditionally heteroscedastic regression errors. However, the techniques to enforce the mean zero moment restriction can also be adapted easily to relax the strong assumption of multivariate normality of the regression errors made in both Hoff and Niu (2012) and Fox and Dunson (2013).

The question of consistency of Bayesian procedures is intimately related to the flexibility of the priors. For instance, in ordinary density estimation problems inclusion of the true density in the KL support of the prior is a sufficient condition to ensure weak consistency via the Schwartz theorem. In density deconvolution problems such a condition is not sufficient but is still required. Theorem 1 from Section 3 thus provides a crucial first step in that direction. We have not pursued the question of consistency of the proposed deconvolution methods any further in this article. We are currently studying the convergence properties of

Bayesian deconvolution models in the simpler known measurement error density set up as the subject of a separate manuscript.

Supplementary Materials

The Supplementary Materials present our arguments in favor of finite mixture models, pointing out how their close connections and their subtle differences with possible infinite dimensional alternatives are exploited to achieve significant reduction in computational complexity while retaining the major advantages of infinite dimensional mixture models including model flexibility and automated model selection and model averaging. The Supplementary Materials also present discussions on the contrasts between regression and measurement errors that preclude the use of covariance regression techniques to model conditionally heteroscedastic measurement errors, the proofs of the theoretical results presented in Section 3, some additional figures, and results of additional simulation experiments. The EATS data analyzed in Section 8 can be accessed from National Cancer Institute arranging a Material Transfer Agreement. R programs implementing our methods are available from the first author upon request.

Acknowledgments

Pati's research was supported by Award No. N00014-14-1-0186 from the Office of Naval Research. Carroll's research was supported in part by a grant R37-CA057030 from the National Cancer Institute. Mallick's research was supported in part by NSF grant DMS0914951. We also acknowledge the Texas A&M University Brazos HPC cluster that contributed to the research reported here. We thank Dr. Anirban Bhattacharya for helpful suggestions.

References

Bhattacharya, A. and Dunson, D. B. (2011). Sparse Bayesian infinite factor models. *Biometrika*, 98, 291-306.

Bhattacharya, A., Pati, D., Pillai, N. and Dunson, D. B. (2014). Bayesian shrinkage. *Unpublished manuscript*.

Brown, P. J. and Griffin, J. E. (2010). Inference with normal-gamma prior distributions in regression problems. *Bayesian Analysis*, 5, 171-188.

Bovy, J., Hogg, D. W. and Rowies, S. T. (2011). Extreme deconvolution: inferring complete distribution functions from noisy, heterogeneous and incomplete observations. *Annals of Applied Statistics*, 5, 1657-1677.

Buonaccorsi, J. P. (2010). *Measurement Error: Models, Methods and Applications*. New York: *Chapman and Hall/CRC*.

Carroll, R. J. and Hall, P. (1988). Optimal rates of convergence for deconvolving a density. *Journal of the American Statistical Association*, 83, 1184-1186.

Carroll, R. J. and Hall, P. (2004). Low order approximations in deconvolution and regression with errors in variables. *Journal of the Royal Statistical Society, Series B*, 66, 31-46.

Carroll, R. J., Ruppert, D., Stefanski, L. A. and Crainiceanu, C. M. (2006). *Measurement Error in Nonlinear Models* (2nd ed.). Boca Raton: *Chapman and Hall/CRC Press*.

Carvalho, M. C., Polson, N. G. and Scott, J. G. (2010). The horseshoe estimator for sparse signals. *Biometrika*, 97, 465-480.

Comte, F. and Lacour, C. (2013). Anisotropic adaptive density deconvolution. *Annales de l'Institut Henri Poincaré - Probabilités et Statistiques*, 49, 569-609.

Devroye, L. (1989). Consistent deconvolution in density estimation. *Canadian Journal of Statistics*, 17, 235-239.

Diggle, P. J. and Hall, P. (1993). A Fourier approach to nonparametric deconvolution of a density estimate. *Journal of the Royal Statistical Society, Series B*, 55, 523-531.

Eilers, P. H. C. and Marx, B. D. (1996). Flexible smoothing with B-splines and penalties. *Statistical Science*, 11, 89-121.

Fan, J. (1991a). On the optimal rates of convergence for nonparametric deconvolution problems. *Annals of Statistics*, 19, 1257-1272.

Fan, J. (1991b). Global behavior of deconvolution kernel estimators. *Statistica Sinica*, 1, 541-551.

Fan, J. (1992). Deconvolution with supersmooth distributions. *Canadian Journal of Statistics*, 20, 155-169.

Fokoué, E. and Titterington, D. M. (2003). Mixtures of factor analyzers. Bayesian estimation and inference by stochastic simulation. *Machine Learning*, 50, 73-94.

Fox, E. B. and Dunson, D. (2013). Bayesian nonparametric covariance regression. *Unpublished manuscript*.

Fraley, C. and Raftery, A. E. (2007). Model-based methods of classification: using the mclust software in chemometrics. *Journal of Statistical Software*, 18, 1-13.

Frühwirth-Schnatter, S. (2006). *Finite Mixture and Markov Switching Models*. New York: *Springer*.

Geweke, J. (2007). Interpretation and inference in mixture models: Simple MCMC works. *Computational Statistics & Data Analysis*, 51, 3529-3550.

Hesse, C. H. (1999). Data driven deconvolution. *Journal of Nonparametric Statistics*, 10, 343-373.

Hoff, P. D. and Niu, X. (2012). A covariance regression model. *Statistica Sinica*, 22, 729-753.

Li, T. and Vuong, Q. (1998). Nonparametric estimation of the measurement error model using multiple indicators. *Journal of Multivariate Analysis*, 65, 139-165.

Liu, M. C. and Taylor, R. L. (1989). A consistent nonparametric density estimator for the

deconvolution problem. *Canadian Journal of Statistics*, 17, 427-438.

Masry, E. (1991). Multivariate probability density deconvolution for stationary random processes. *IEEE Transactions on Information Theory*, 37, 1105-1115.

Mengersen, K. L., Robert, C. P. and Titterington, D. M. (eds) (2011). *Mixtures - Estimation and Applications*. Chichester: *John Wiley*.

Neumann, M. H. (1997). On the effect of estimating the error density in nonparametric deconvolution. *Journal of Nonparametric Statistics*, 7, 307-330.

Pelenis, J. (2014). Bayesian regression with heteroscedastic error density and parametric mean function. *Journal of Econometrics*, 178, 624-638.

Rousseau, J. and Mengersen, K. (2011). Asymptotic behavior of the posterior distribution in overfitted mixture models *Journal of the Royal Statistical Society, Series B*, 73, 689-710.

Sarkar, A., Mallick, B. K., Staudenmayer, J., Pati, D. and Carroll, R. J. (2014). Bayesian semiparametric density deconvolution in the presence of conditionally heteroscedastic measurement errors. Forthcoming in *Journal of Computational and Graphical Statistics*.

Staudenmayer, J., Ruppert, D. and Buonaccorsi, J. P. (2008). Density estimation in the presence of heteroscedastic measurement error. *Journal of the American Statistical Association*, 103, 726-736.

Subar, A. F., Thompson, F. E., Kipnis, V., Midthune, D., Hurwitz, P. McNutt, S., McIntosh, A. and Rosenfeld, S. (2001). Comparative validation of the block, Willet, and National Cancer Institute food frequency questionnaires. *American Journal of Epidemiology*, 154, 1089-1099.

Youndjé, E. and Wells, M. T. (2008). Optimal bandwidth selection for multivariate kernel deconvolution. *TEST*, 17, 138-162.

True Error Distribution	Covariance Structure	Sample Size	MISE $\times 10^4$		
			MLFA	MIW	Naive
(a) Multivariate Normal	I	500	**1.24**	3.05	8.01
		1000	**0.59**	1.33	6.58
	LF	500	6.88	**6.33**	33.41
		1000	5.15	**3.10**	32.42
	AR	500	11.91	**5.51**	27.17
		1000	9.82	**2.78**	26.01
	EXP	500	7.15	**4.40**	17.82
		1000	5.46	**2.19**	17.40
(b) Mixture of Multivariate Normal	I	500	**1.28**	3.24	5.97
		1000	**0.64**	1.37	4.99
	LF	500	**7.28**	7.51	31.62
		1000	**4.17**	4.34	31.48
	AR	500	10.43	**6.66**	30.74
		1000	7.75	**4.35**	28.90
	EXP	500	7.16	**5.18**	17.85
		1000	4.87	**2.66**	17.26

Table 1: Mean integrated squared error (MISE) performance of MLFA (mixtures of latent factor analyzers) and MIW (mixtures with inverse Wishart priors) density deconvolution models described in Section 2 of this article for **homoscedastic** errors compared with a naive method that ignores measurement errors for different measurement error distributions. The minimum value in each row is highlighted.

True Error Distribution	Covariance Structure	Sample Size	MISE $\times 10^4$		
			MLFA	MIW	Naive
(a) Multivariate Normal	I	500	**2.53**	19.08	10.64
		1000	**1.15**	9.43	9.14
	LF	500	**11.46**	34.21	21.33
		1000	**5.78**	15.98	20.75
	AR	500	**17.11**	30.83	36.44
		1000	**10.77**	12.46	36.37
	EXP	500	**11.63**	26.99	24.28
		1000	**6.67**	10.56	23.36
(b) Mixture of Multivariate Normal	I	500	**2.79**	22.17	20.16
		1000	**1.38**	10.55	19.39
	LF	500	**13.39**	35.67	43.43
		1000	**7.50**	20.86	43.28
	AR	500	**18.27**	35.70	75.26
		1000	**12.06**	16.64	77.55
	EXP	500	**12.11**	34.50	48.76
		1000	**7.59**	13.74	50.02

Table 2: Mean integrated squared error (MISE) performance of MLFA (mixtures of latent factor analyzers) and MIW (mixtures with inverse Wishart priors) density deconvolution models described in Section 2 of this article for **conditionally heteroscedastic** errors compared with a naive method that ignores measurement errors or different measurement error distributions. The minimum value in each row is highlighted.

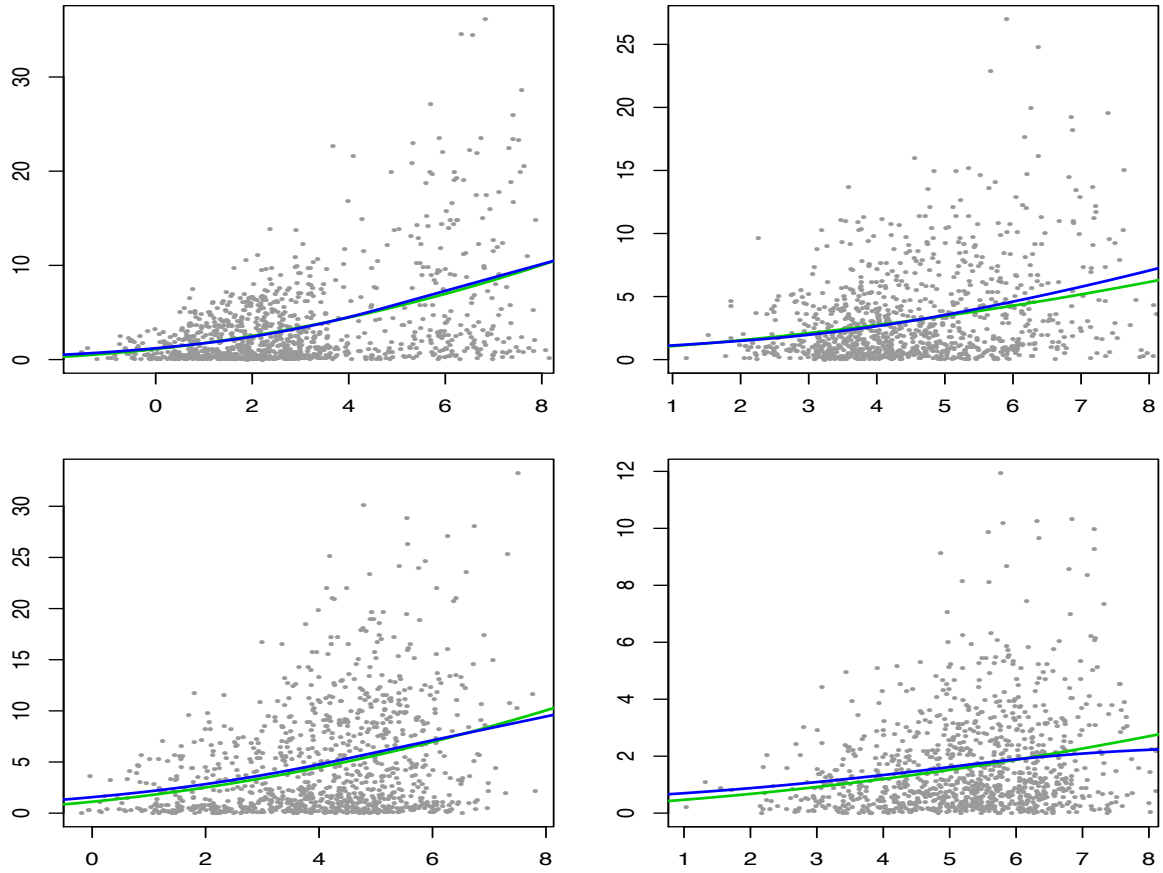

Figure 1: Results for the variance functions $s^2(X)$ produced by the univariate density deconvolution method for each component of $\mathbf{X}$ for the conditionally heteroscedastic error distribution $f_{\boldsymbol{\epsilon}}^{(2)}$ with sample size $n = 1000$, $m_i = 3$ replicates for each subject and identity matrix (I) for the component specific covariance matrices. The results correspond to the data set that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets for the MLFA (mixtures of latent factor analyzers) method. For each component of $\mathbf{X}$, the true variance function is $s^2(X) = (1 + X/4)^2$. See Section 2.2.2 and Section 6 for additional details. In each panel, the true (lighter shaded lines) and the estimated (darker shaded lines) variance functions are superimposed over a plot of subject specific sample means vs subject specific sample variances. The figure is in color in the electronic version of this article.

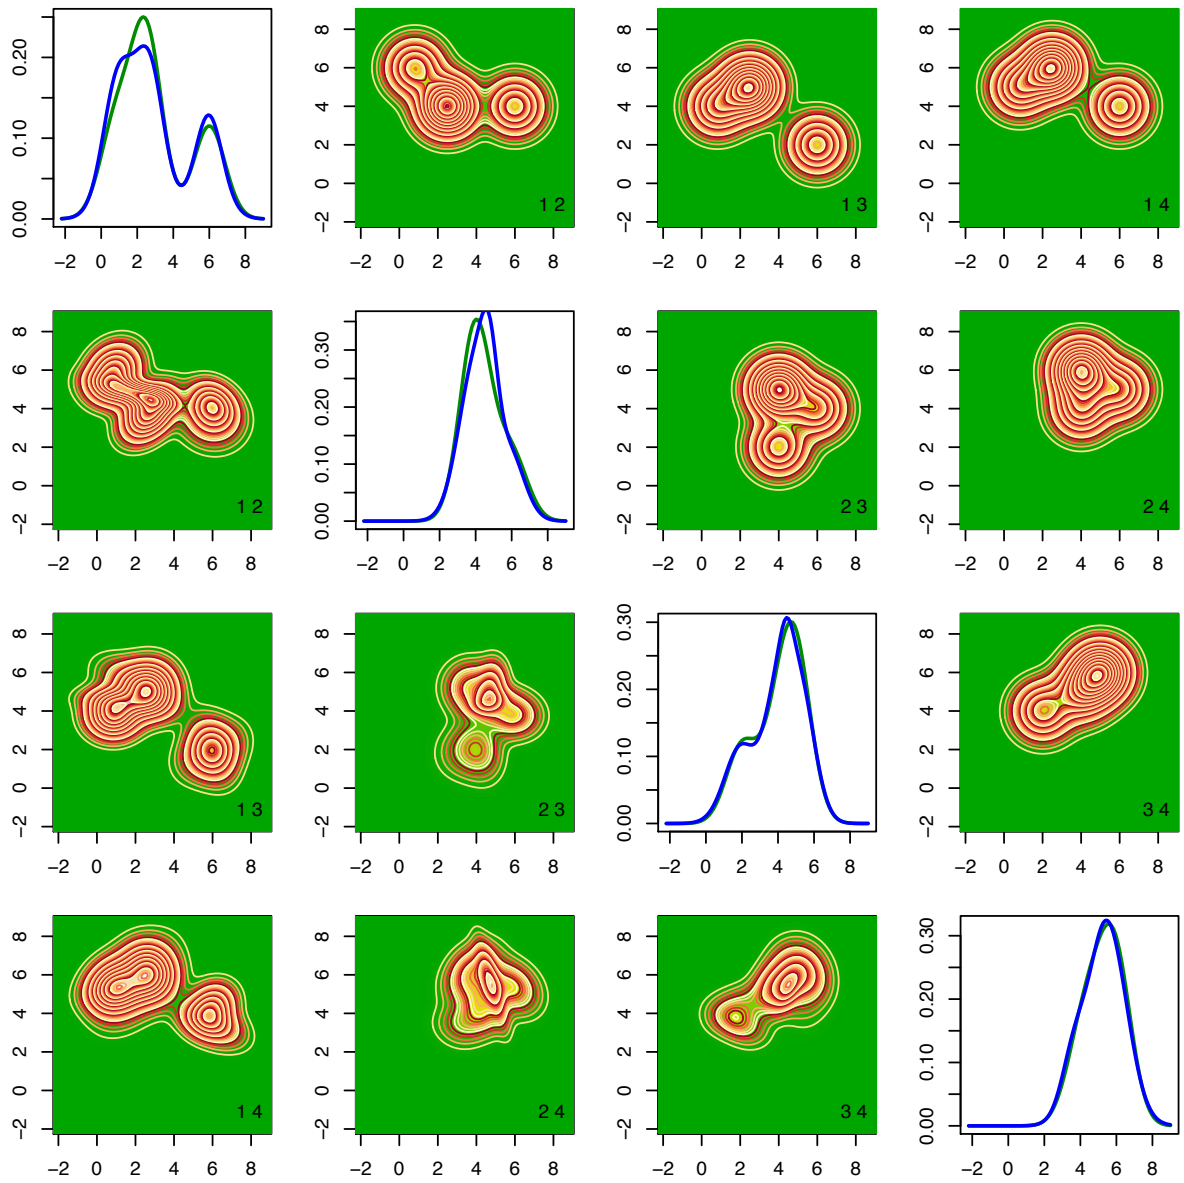

Figure 2: Results for $f_{\mathbf{X}}$ produced by the MIW (mixtures with inverse Wishart priors) method for the conditionally heteroscedastic error distribution $f_{\boldsymbol{\epsilon}}^{(2)}$ with sample size $n = 1000$, $m_i = 3$ replicates for each subject and identity matrix (I) for the component specific covariance matrices. The results correspond to the data set that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets. See Section 7 for additional details. The upper triangular panels show the contour plots of the true two dimensional marginal densities. The lower triangular diagonally opposite panels show the corresponding estimates. The diagonal panels show the true (lighter shaded lines) and the estimated (darker shaded lines) one dimensional marginals. The figure is in color in the electronic version of this article.

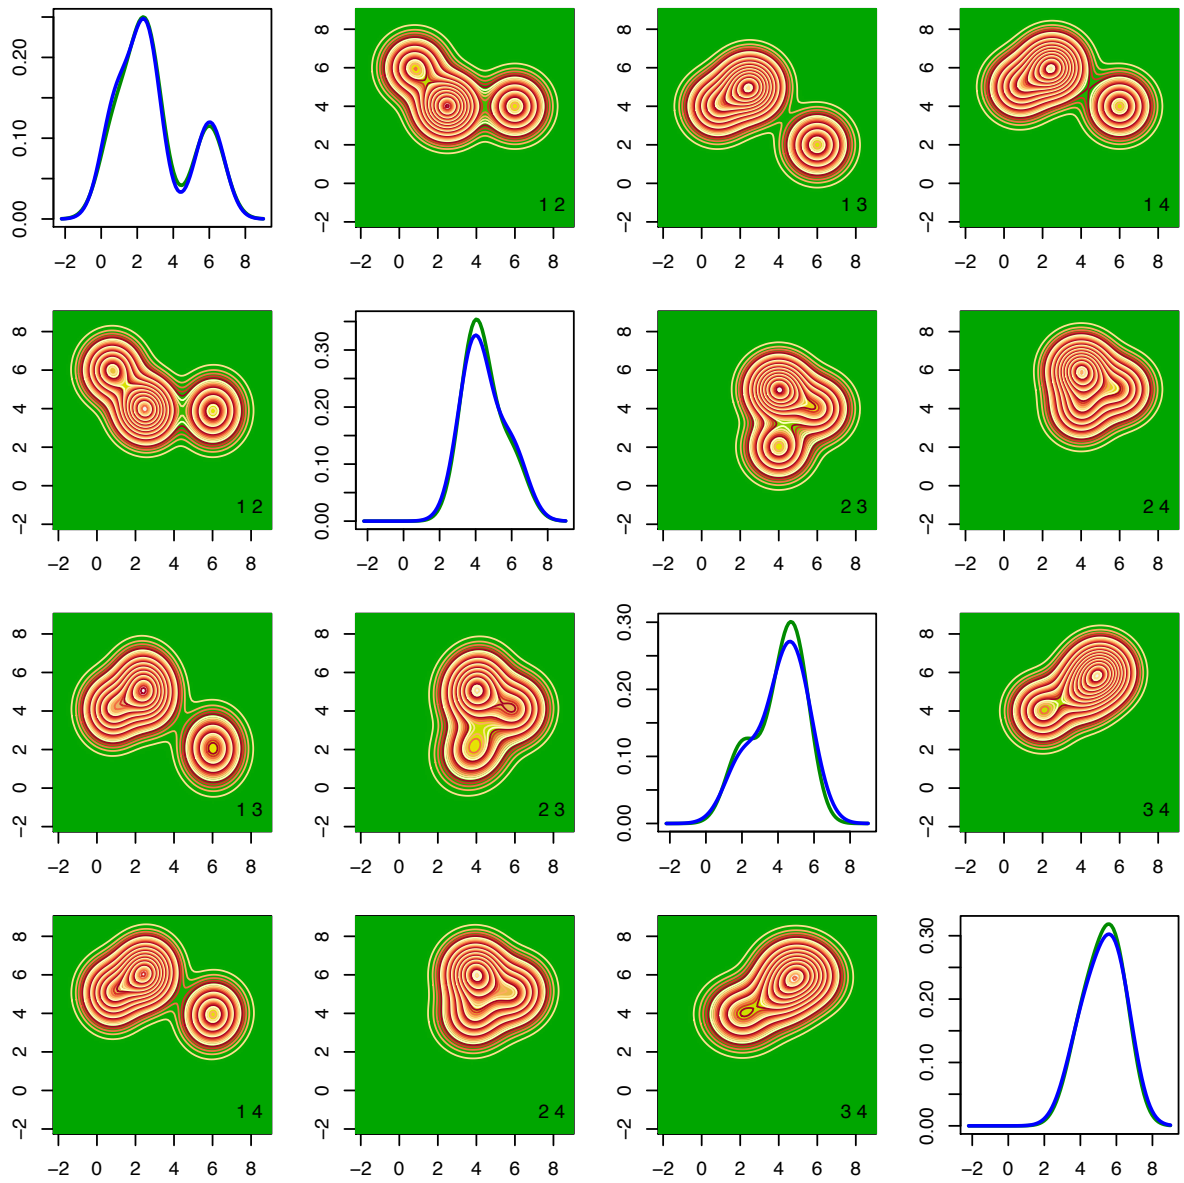

Figure 3: Results for the $f_{\mathbf{X}}$ produced by the MLFA (mixtures of latent factor analyzers) method for the conditionally heteroscedastic error distribution $f_{\boldsymbol{\epsilon}}^{(2)}$ with sample size $n = 1000$, $m_i = 3$ replicates for each subject and identity matrix (I) for the component specific covariance matrices. The results correspond to the data set that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets. See Section 7 for additional details. The upper triangular panels show the contour plots of the true two dimensional marginal densities. The lower triangular diagonally opposite panels show the corresponding estimates. The diagonal panels show the true (lighter shaded lines) and the estimated (darker shaded lines) one dimensional marginals. The figure is in color in the electronic version of this article.

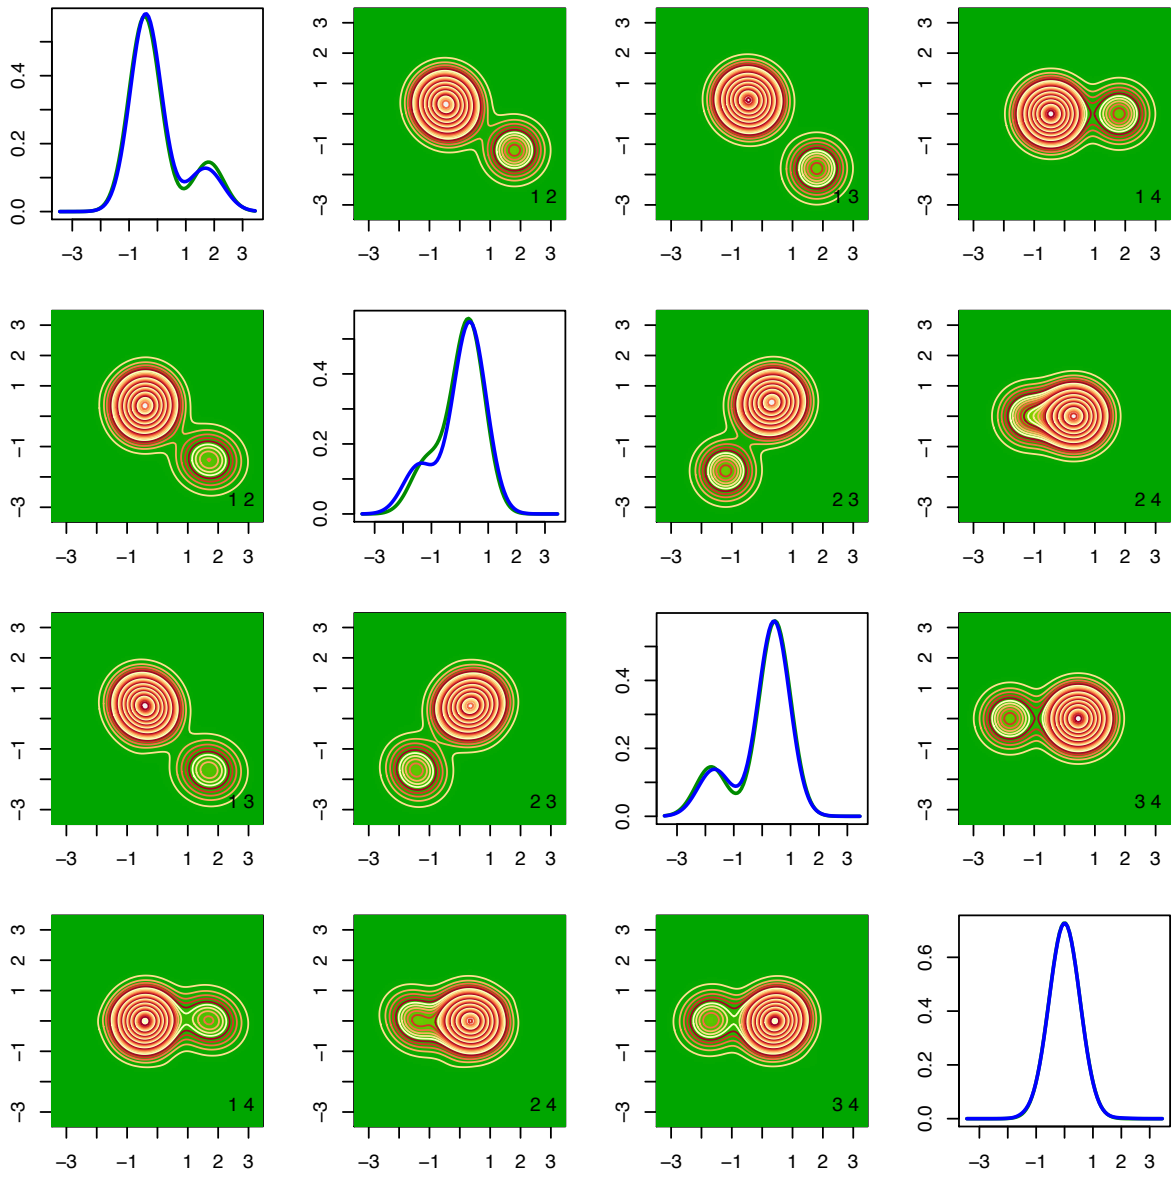

Figure 4: Results for the density of the scaled measurement errors $f_{\boldsymbol{\epsilon}}$ produced by the MIW (mixtures with inverse Wishart priors) method for the conditionally heteroscedastic error distribution $f_{\boldsymbol{\epsilon}}^{(2)}$ with sample size $n = 1000$, $m_i = 3$ replicates for each subject and identity matrix (I) for the component specific covariance matrices. The results correspond to the data set that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets. See Section 7 for additional details. The upper triangular panels show the contour plots of the true two dimensional marginal densities. The lower triangular diagonally opposite panels show the corresponding estimates. The diagonal panels show the true (lighter shaded lines) and the estimated (darker shaded lines) one dimensional marginals. The figure is in color in the electronic version of this article.

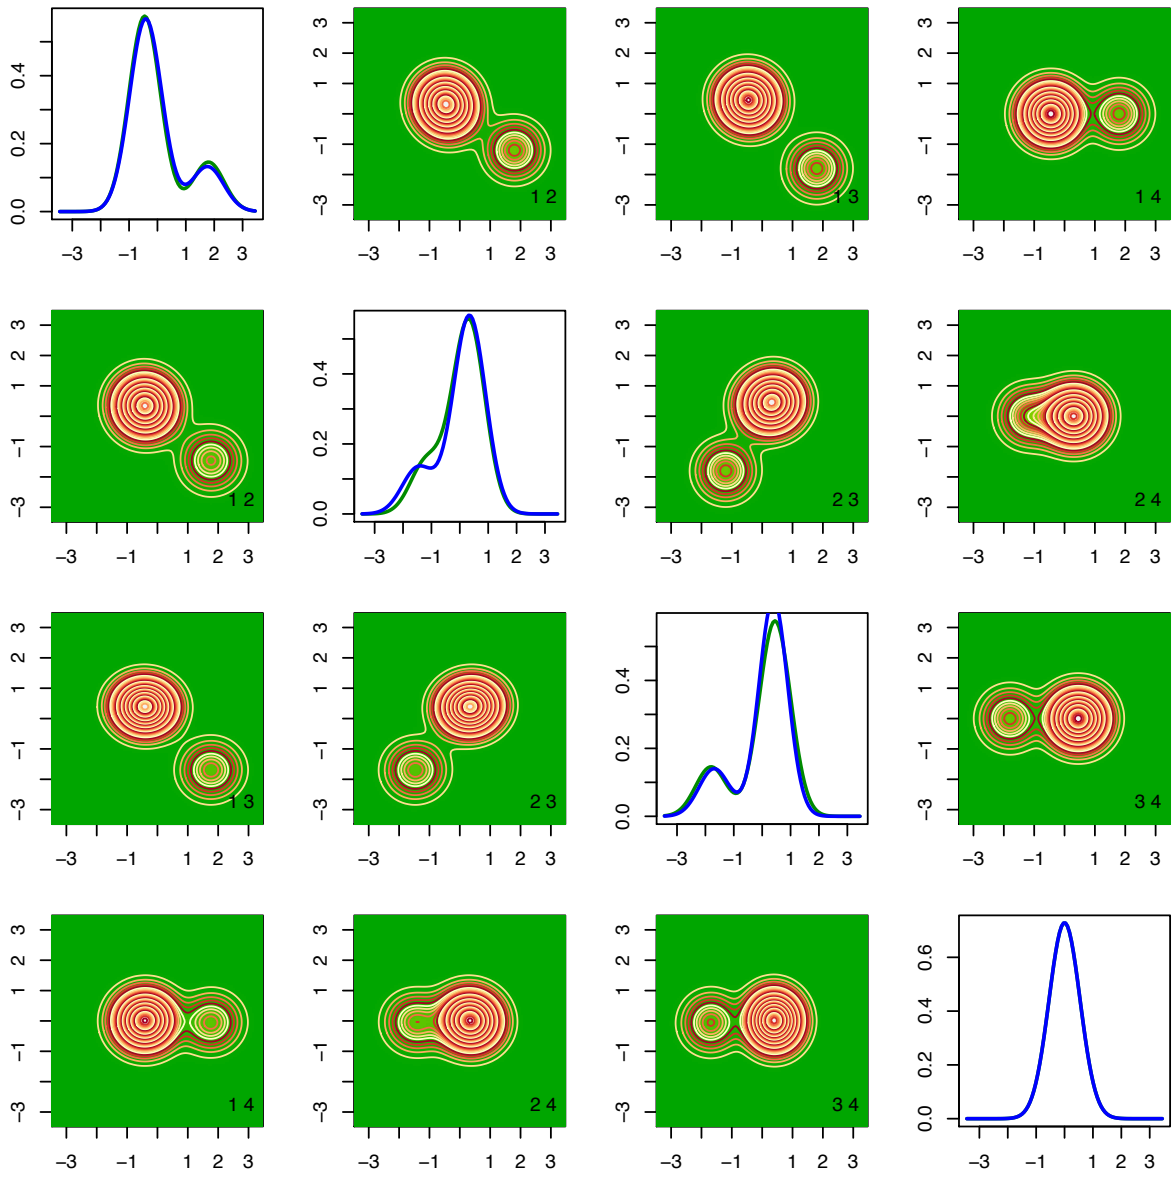

Figure 5: Results for the density of the scaled measurement errors $f_{\boldsymbol{\epsilon}}$ produced by the MLFA (mixtures of latent factor analyzers) method for the conditionally heteroscedastic error distribution $f_{\boldsymbol{\epsilon}}^{(2)}$ with sample size $n = 1000$, $m_i = 3$ replicates for each subject and identity matrix (I) for the component specific covariance matrices. The results correspond to the data set that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets. See Section 7 for additional details. The upper triangular panels show the contour plots of the true two dimensional marginal densities. The lower triangular diagonally opposite panels show the corresponding estimates. The diagonal panels show the true (lighter shaded lines) and the estimated (darker shaded lines) one dimensional marginals. The figure is in color in the electronic version of this article.

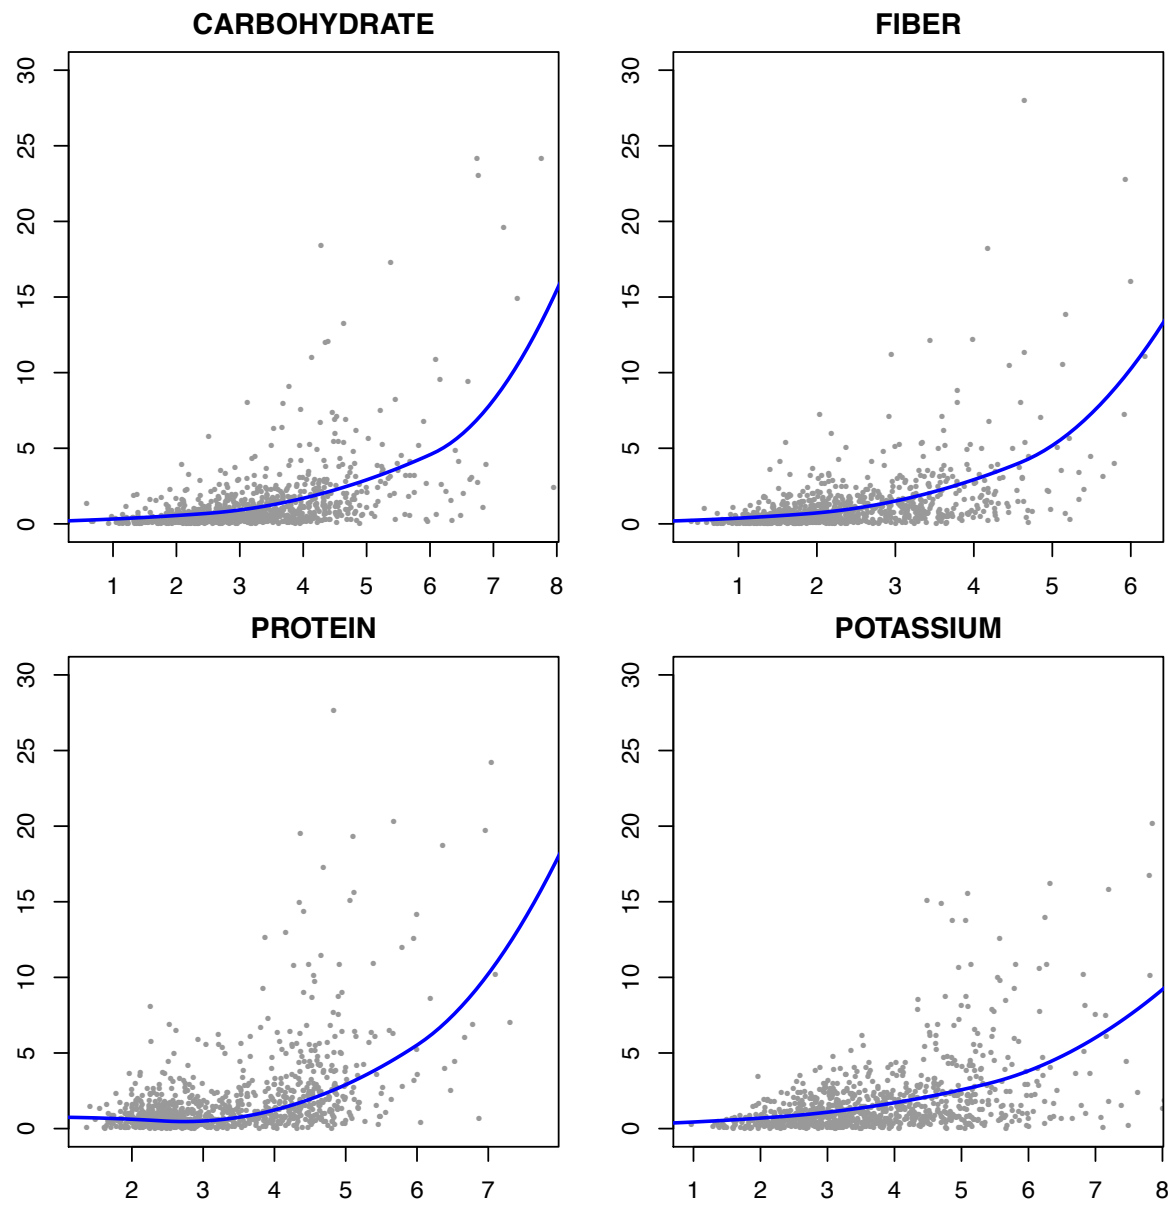

Figure 6: Estimated variance functions $s^2(X)$ produced by the univariate density deconvolution method for each component of $\mathbf{X}$ for the EATS data set with sample size $n = 965$, $m_i = 4$ replicates for each subject. See Section 8 for additional details. The figure is in color in the electronic version of this article.

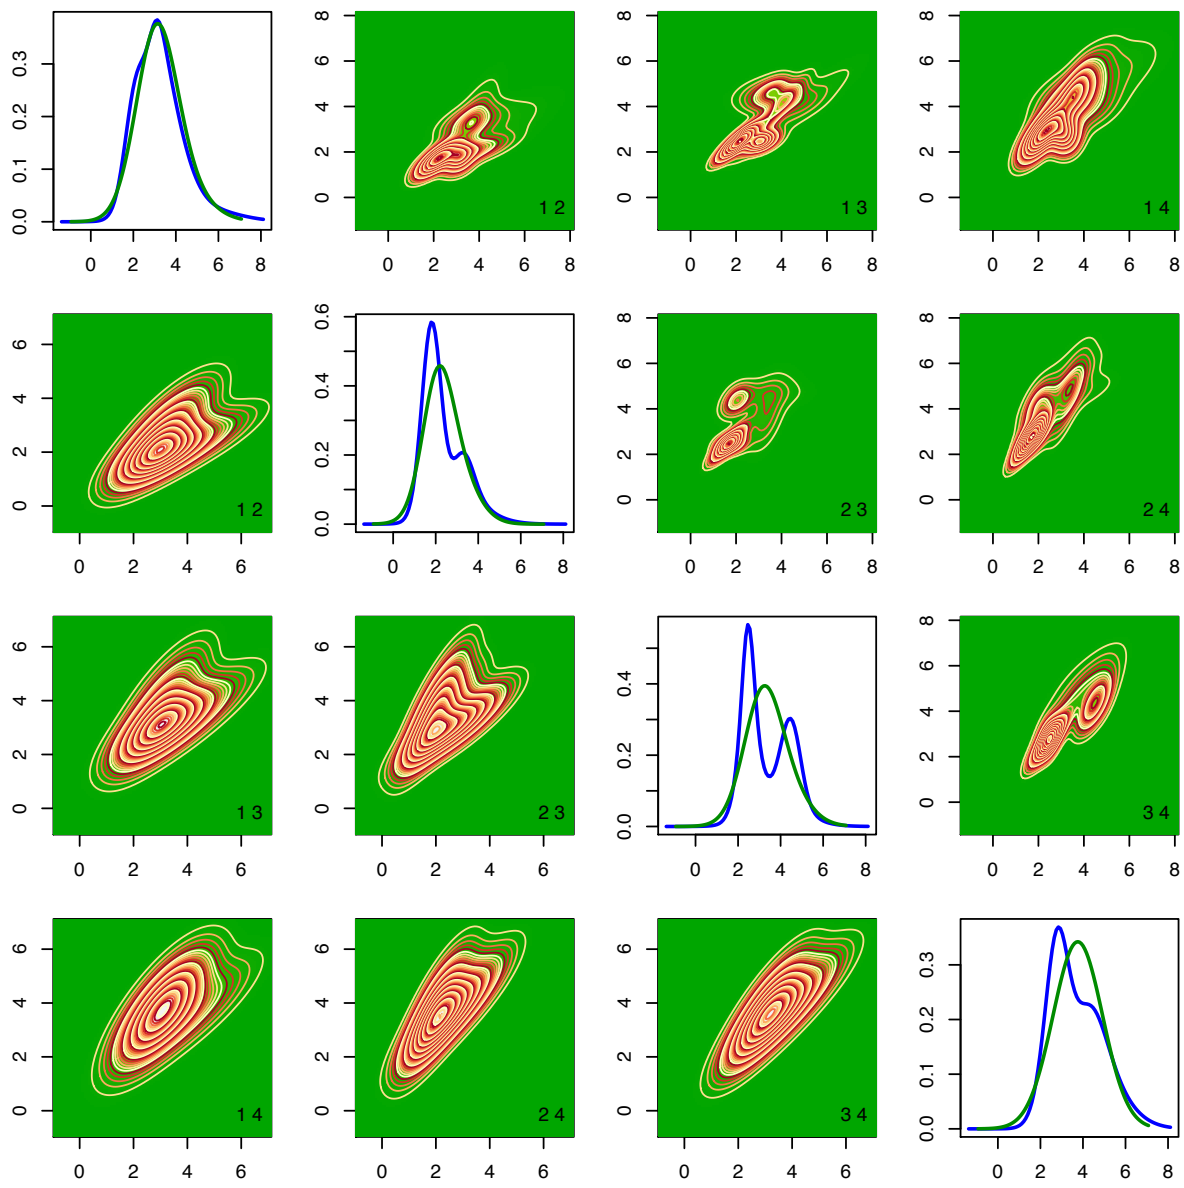

Figure 7: Results for the EATS data set for the $f_{\mathbf{X}}$. The off-diagonal panels show the contour plots of two-dimensional marginals estimated by the MIW method (upper triangular panels) and the MLFA method (lower triangular panels). The diagonal panels show the one dimensional marginal densities estimated by the MIW method (darker shaded lines) and the MLFA method (lighter shaded lines). The figure is in color in the electronic version of this article.

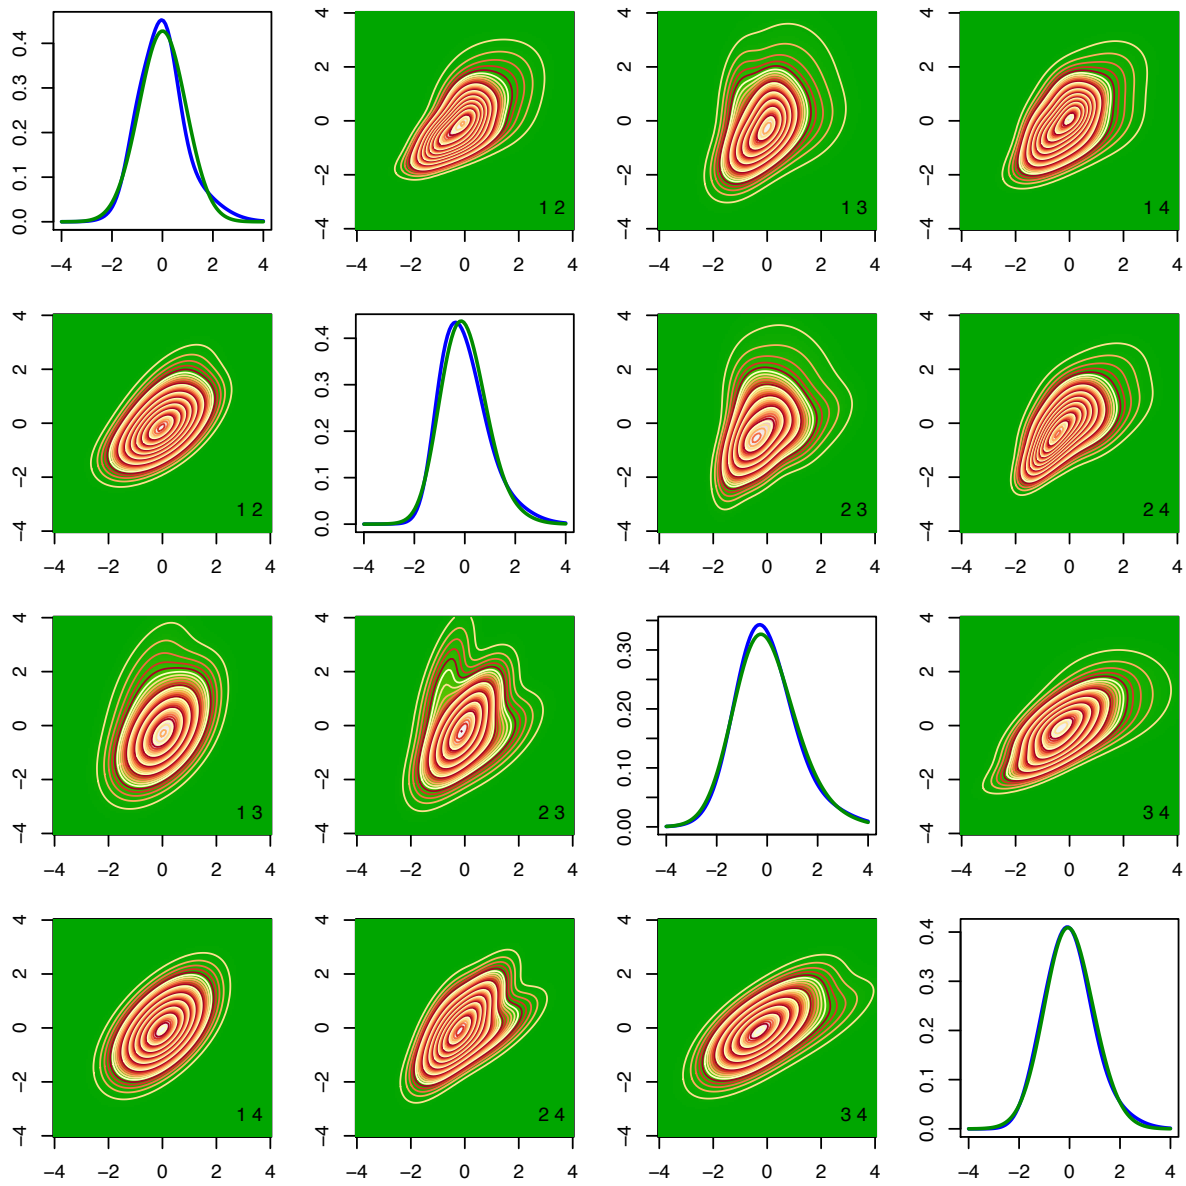

Figure 8: Results for the EATS data set for the density of the scaled errors $f_{\boldsymbol{\epsilon}}$. The off-diagonal panels show the contour plots of two-dimensional marginals estimated by the MIW method (upper triangular panels) and the MLFA method (lower triangular panels). The diagonal panels show the one dimensional marginal densities estimated by the MIW method (darker shaded lines) and the MLFA method (lighter shaded lines). The figure is in color in the electronic version of this article.

Supplementary Materials for **Bayesian Semiparametric Multivariate Density Deconvolution**

Abhra Sarkar
Department of Statistics, Texas A&M University, 3143 TAMU, College Station, TX 77843-3143 USA
abhra@stat.tamu.edu

Debdeep Pati
Department of Statistics, Florida State University, Tallahassee, FL 32306-4330 USA
debdeep@stat.fsu.edu

Bani K. Mallick
Department of Statistics, Texas A&M University, 3143 TAMU, College Station, TX 77843-3143 USA
bmallick@stat.tamu.edu

Raymond J. Carroll
Department of Statistics, Texas A&M University, 3143 TAMU, College Station, TX 77843-3143 USA
carroll@stat.tamu.edu

S.1 Finite Mixture Models vs Infinite Mixture Models

In this article, we modeled the $f_{\mathbf{X}}$ and the density of the scaled measurement errors $f_{\boldsymbol{\epsilon}}$ using mixtures of fixed finite number of multivariate normal kernels. Alternative approaches that escape the need to prespecify the number of mixture components include models with potentially infinite number of mixture components, models induced by Dirichlet processes (Ferguson, 1973; Escobar and West, 1995) being perhaps the most popular among such techniques. Apart from flexibility, one major advantage of such techniques comes from the ability of associated MCMC machinery to perform model selection and model averaging implicitly and semiautomatically. Model averaging is achieved by allowing the number of mixture components to vary from one MCMC iteration to the other. The number of mixture components that is visited the maximum number of times by the sampler then provides a maximum a-posteriori (MAP) estimate of the number of mixture components required to approximate the target density. However, in complicated multivariate set up like ours, MCMC algorithms for such infinite dimensional models become computationally highly intensive.

Mixtures based on fixed finite number of components, on the other hand, can greatly reduce computational complexity. Recent studies of asymptotic properties of the posterior of overfitted mixture models (Rousseau and Mengersen, 2011) suggest that mixture models with sufficiently large number of components can perform automatic model selection and model averaging just like infinite dimensional models. Additionally, as the proofs of the results in Section 3 imply, the use of mixture models with fixed finite number of components does not necessarily imply a compromise on the issue of flexibility. The approaches adopted in this article try to take the best from both worlds. Computational burden is reduced by keeping the number of mixture components fixed at some finite values. At the same time, simultaneous semiautomatic model selection and model averaging is achieved by exploiting asymptotic properties of overfitted mixture models. We elaborate our arguments below, pointing out the close connections and the subtle differences our adopted finite dimensional models have with the aforementioned infinite dimensional alternatives.

S.1.1 Computational Complexity

The finite mixtures of multivariate normal kernels with symmetric Dirichlet priors that we used in this article to model both the $f_{\mathbf{X}}$ and the density of the scaled measurement errors $f_{\boldsymbol{\epsilon}}$ have close connections with infinite dimensional Dirichlet process based mixture models. Indeed, as $K_{\mathbf{X}} \to \infty$, the finite dimensional symmetric Dirichlet prior (3) used to model $f_{\mathbf{X}}$ converges weakly to a countably infinite distribution distributed according to a Dirichlet process with concentration parameter $\alpha_{\mathbf{X}}$ (Ishwaran and Zarepour, 2000, 2002). Our proposed mechanism to enforce the mean zero restriction on $f_{\boldsymbol{\epsilon}}$ specifically requires a finite dimensional symmetric prior on the mixture probabilities and therefore does not admit a straightforward infinite dimensional extension. But in the limit, as $K_{\boldsymbol{\epsilon}} \to \infty$, a reformulation of the model results in a complicated multivariate version of the infinite dimensional model of Sarkar, et al. (2014) (See Lemma 4 in Section 6). The implementation of such complex models, specially the complicated mean restricted model for the scaled errors, will be computationally intensive in a multivariate setting like ours. The computational simplicity of the finite dimensional methods proposed in this article make them particularly suitable for multivariate problems.

In this paragraph, we discuss additional mixing issues that render infinite dimensional models, particularly the ones with non or semiconjugate priors on the component specific parameters (like our MLFA model), unsuitable for multivariate applications. There are two main types of MCMC algorithms for fitting infinite dimensional mixture models - conditional methods and marginal methods. In the conditional scheme, the mixture probabilities are sampled. The mixture labels are then updated independently, conditional on the mixture probabilities. The mixture probabilities in infinite dimensional mixture models can be

stochastically ordered. For instance, mixture probabilities in a Dirichlet process mixture model satisfy $E(\pi_k) > E(\pi_{k+1})$ and $\Pr(\pi_k > \pi_{k+1}) > 0.5$ for all $k \in \mathbb{N}$. This imposes weak identifiability on the mixture labels resulting in a complicated model space comprising many local modes of varying importance. Different permutations of the mixture labels are not equivalent and exploration of the entire model space becomes important for valid inference. In high dimensional and large data settings this is difficult to achieve even by sophisticated MCMC algorithms with carefully designed label switching moves (Hastie, et al., 2013). The problem can be avoided with marginal methods (Neal, 2000) that integrate out the mixture probabilities and work with the resulting Polya urn scheme, rendering the mixture labels dependent but nonidentifiable. Unfortunately, such integration is possible only when conjugate priors are assigned to the component specific parameters. Typically for infinite dimensional models with non or semiconjugate priors on the component specific parameters, good mixing is thus difficult to achieve, particularly in complicated multivariate setup like ours. On the contrary, the issues of mixing and convergence become much less important for finite mixture models with exchangeable priors on the mixture probabilities. With $K_{\mathbf{X}}$ and $K_{\boldsymbol{\epsilon}}$ mixture components for the densities $f_{\mathbf{X}}$ and $f_{\boldsymbol{\epsilon}}$, respectively, the posterior is still multimodal but comprises $K_{\mathbf{X}}! \times K_{\boldsymbol{\epsilon}}!$ modal regions that are exact copies of each other. For inference on the overall density or any other functions of interest that are invariant to permutations of the mixture labels, it is only important that the MCMC sampler visits and explores at least one of the modal regions well and label switching (or the lack of it) does not present any problem (Geweke, 2007).

S.1.2 Model Selection and Model Averaging

As mentioned at the beginning of Section S.1, a major advantage of infinite dimensional mixture models is their ability to implicitly and semiautomatically perform model selection and model averaging. Properties of overfitted mixture models can be exploited to achieve the same in finite dimensional models with sufficiently large number of components. Recently Rousseau and Mengersen (2011) studied the asymptotic behavior of the posterior for overfitted mixture models with Dirichlet prior $\text{Dir}(\alpha_1, \ldots, \alpha_K)$ on the mixture probabilities in a measurement error free set up and showed that the hyper parameter $(\alpha_1, \ldots, \alpha_k)$ strongly influences the way the posterior handles overfitting. In particular, when $\max_{k=1,\ldots,K} \alpha_k < L/2$, where L denotes the number of parameters specifying the component kernels, the posterior is asymptotically stable and concentrates in regions with empty redundant components. In this article, we chose symmetric Dirichlet priors $\text{Dir}(\alpha/K, \ldots, \alpha/K)$ on the mixture probabilities to model both the $f_{\mathbf{X}}$ and the density of the scaled measurement errors $f_{\boldsymbol{\epsilon}}$. We set $\alpha_{\mathbf{X}} = \alpha_{\boldsymbol{\epsilon}} = 1$ so that the condition $\alpha/K < L/2$ is satisfied for both $f_{\mathbf{X}}$ and $f_{\boldsymbol{\epsilon}}$. In simulation experiments reported in Section 7, the behavior of the posterior was similar to that observed

by Rousseau and Mengersen (2011) in measurement error free set up. That is, when $K_{\mathbf{X}}$ and $K_{\boldsymbol{\epsilon}}$ were assigned sufficiently large values, the MCMC chain quickly reached a stable stage where the redundant components became empty. See Figure S.3 in the main article and Figure S.9 and S.10 in the Supplementary Materials for illustrations, where, with some abuse of nomenclature, the k^{th} component is called empty if the associated mixture probability $\pi_k \leq 0.05$. Since such overfitted mixture models allow the number of nonempty mixture components to vary from one MCMC iteration to the next, model averaging is automatically achieved. MAP estimates of the numbers of mixture components required to approximate the target densities are given by the numbers of components which are visited the maximum number of times by the MCMC sampler, as in the case of infinite mixture models.

As discussed in the main paper, for the MIW method, when the measurement errors are conditionally heteroscedastic and the true covariance matrices are highly sparse, the strategy usually overestimates the number of non-empty mixture components required to approximate the target densities well. In these cases, the MIW method becomes highly numerically unstable and much larger sample sizes are required for the asymptotic results to hold. See Figure S.2 in the main article for an illustration. This may be regarded more as a limitation of the MIW method than a limitation of the adopted strategy to determine $K_{\mathbf{X}}$ and $K_{\boldsymbol{\epsilon}}$. For the numerically more stable MLFA model, the asymptotic results are valid even for moderate sample sizes and such models are also more robust to overestimation of the number of nonempty clusters.

S.1.3 Model Flexibility

The proofs of the support results presented in Section 3 require that the number of mixture components of the corresponding mixture models be allowed to vary over the set of all positive integers. However, as the technical details of the proofs reveal, the use of mixture models with fixed finite number of components does not necessarily imply a compromise on the issue of flexibility. Indeed, a common recurring idea in the proofs of all these results, including those for the variance functions, is to show that any function coming from the target class can be approximated with any desired level of accuracy by the corresponding finite mixture models provided the models comprise sufficiently large number of mixture components and the function satisfies some fairly minimal regularly conditions. The requirement that the priors on the number of mixture components assign positive probability to all positive integers only helps us reach the final conclusions as immediate consequences. For any given data set of finite size, the number of mixture components required to approximate a target density will always be bounded above by the number of latent or observed variables generated by the target density. For most practical applications the required number would actually be much smaller than the number of variables generated by the target. Even if one applies mixture

models that a-priori allow potentially infinitely many mixture components, the posterior will essentially concentrate on a finite set comprising moderately small positive integers. This means that for all practical purposes solutions based on finite mixture models with fixed but sufficiently large number of mixture components will essentially be as robust as solutions based on their infinite or varying dimensional counterparts while at the same time being significantly less burdensome from a computational viewpoint. The requirement that the priors on the number of mixture components assign positive mass on *all* positive integers may thus be relegated to the requirement that the priors assign positive mass on sets of the form $\{1, \ldots, K\}$, where K is sufficiently large. Posterior computation for such models might be even much more intensive and complex requiring reversible jump moves. Since a mixture model with K components is at least as flexible as a model with $(K-1)$ components, properties of overfitted mixture models discussed in Section S.1.2 allow us to adopt a much simpler strategy. We can simply keep the number of mixture components fixed at sufficiently large values for all MCMC iterations. Carefully chosen priors for the mixture probabilities then result in a posterior that concentrates in regions favoring empty redundant components, essentially eliminating the need to assign any priors on the number of mixture components. We will still need some mechanism, preferably an automated and data adaptive one, to determine what values of K would be sufficiently large. This issue is discussed in the section on hyper-parameter choices in Section 4.

The discussions of Section S.1 suggest that finite mixture models with sufficiently large number of mixture components and carefully chosen priors for the mixture probabilities can essentially retain the major advantages of infinite dimensional alternatives including flexibility, automated model averaging and model selection while at the same time being computationally much less burdensome, making them our preferred choice for complicated high dimensional problems.

S.2 Regression Errors vs Measurement Errors

Consider the problem of flexible modeling of conditionally heteroscedastic regression errors where the response and the covariates are both univariate. Consider also the problem of modeling conditionally heteroscedastic measurement errors in a univariate deconvolution set up. From a modeling perspective, Bayesian hierarchical framework allows us to treat these two problems on par by treating both the covariate in the regression problem and the variable of interest in the deconvolution problem simply as conditioning variables. Of course in the regression problem X is precisely measured, whereas in the deconvolution problem X would be latent, but in either case we are required to flexibly model the density of $(U \mid X)$ subject to $E(U \mid X) = 0$, where U, depending upon the context, denotes either regression or measurement errors. Models for regression errors that allow their variance to vary with the values of the covariate (Pati and Dunson, 2013; Pelenis, 2014) can thus be tried as potential candidates for models for univariate conditionally heteroscedastic measurement errors. Conversely, the models for conditionally heteroscedastic univariate measurement errors (Staudenmayer, et al., 2008; Sarkar, et al., 2014) can also be employed to model univariate conditionally heteroscedastic regression errors.

This is not quite true in a multivariate set up. Interpreting the variables of interest $\mathbf{X}$ broadly as conditioning variables, one can again loosely connect the problem of modeling conditionally heteroscedastic multivariate measurement errors to the problem of covariance regression (Hoff and Niu, 2012; Fox and Dunson, 2013 etc.), where the goal is to develop models that allow the covariance of multivariate regression errors to vary flexibly with precisely measured and possibly multivariate predictors. In covariance regression problems, the dimension of the regression errors is typically unrelated to the dimension of the predictors and different components of the regression errors are assumed to be equally influenced by different components of the predictors. In multivariate deconvolution problems, in contrast, the dimension of each $\mathbf{U}_{ij}$ is exactly the same as the dimension of $\mathbf{X}_i$. Although different components of the measurement error vectors $\mathbf{U}_{ij}$ may be correlated, since the ℓ^{th} component $U_{ij\ell}$ is the measurement error associated exclusively with $X_{i\ell}$, the dependence of $U_{ij\ell}$ on $\mathbf{X}_i$ can be explained mostly through $X_{i\ell}$. Figure 6, for instance, suggests strong conditional heteroscedasticity patterns and it is plausible to assume that the conditional variability in $U_{ij\ell}$ can be modeled as a function of $X_{i\ell}$ only. The characteristics of multivariate measurement errors are, therefore, very different from the characteristics of multivariate regression errors. These contrasts preclude direct application of covariance regression approaches to model conditionally heteroscedastic multivariate measurement errors. The converse is also true. Models for conditionally heteroscedastic multivariate measurement errors should accommodate the aforementioned unique features and should not be used for modeling predictor dependent multivariate regression errors.

S.3 Conditionally Heteroscedastic Covariance Model

In Figure 6 in the main paper, we showed the plots of subject specific means $\overline{W}_{i\ell}$ of the replicates vs the corresponding subject-specific variances $S^2_{W,i\ell}$ for each of the four dietary components included in our analysis in Section 8. These plots suggest very strong conditional heteroscedasticity patterns in the measurement errors. If we consider the plots of subject specific means $\overline{W}_{i\ell}$ vs subject specific variances $S^2_{W,im}$ for all possible pairs (ℓ, m), we will see similar monotone increasing patterns not just for the pairs with $\ell = m$, but in pairs with $\ell \neq m$ too. This can be explained by the high correlation between different components of $\mathbf{X}_i$, see Figure 7, and does not imply that the conditional variability in $U_{ij\ell}$ depends on other components of $\mathbf{X}_i$, not just $X_{i\ell}$. As discussed in Section S.2 of the Supplementary Materials, since the ℓ^{th} component $U_{ij\ell}$ is the measurement error associated exclusively with $X_{i\ell}$, it is plausible to assume that the conditional variability of $U_{ij\ell}$ can be modeled mostly as a function of $X_{i\ell}$ only. We present here some diagnostic plots to further validate this structural assumption. Figure S.1 shows the plots of $\widehat{X}_{i\ell}$ vs subject specific variances $\widehat{S}^2_{\epsilon,im}$ of $\widehat{\epsilon}_{ijm}$, where $\widehat{X}_{i\ell}$ represent the posterior means of $X_{i\ell}$ values and $\widehat{\epsilon}_{ijm} = (W_{ijm} - \widehat{X}_{im})/\widehat{s}_m(\widehat{X}_{im})$ represent the corresponding scaled measurement error residuals produced by the univariate submodels for the EATS data set analyzed in Section 8 of the main paper. The figure indicates constant variance of the scaled measurement error residuals $\widehat{\epsilon}_{ij\ell}$ over the entire range of X_{im} values for all (ℓ, m) pairs. Nonparametric Eubank-Hart tests of no covariate effect (Eubank and Hart, 1992) applied to $(\widehat{X}_{i\ell}, \widehat{S}^2_{\epsilon,im})$ for all (ℓ, m) pairs (treating $\widehat{X}_{i\ell}$ as the covariate and $\widehat{S}^2_{\epsilon,im}$ as the response) produced a minimum p-value of 0.217, strongly suggesting that there is no residual heteroscedasticity left in $U_{ij\ell}$ after accounting for the variability in $U_{ij\ell}$ that can be explained through $X_{i\ell}$ only. Model (16) for conditionally heteroscedastic multivariate measurement errors developed in Section 2.2.2 of the main paper implied $\text{var}(U_{ij\ell} \mid \mathbf{X}_i) = s^2_\ell(X_{i\ell})\text{var}(\epsilon_{ij\ell})$, properly accommodating this unique structural feature of conditionally varying multivariate measurement errors.

Model (16) also assumed that only the conditional variability of $\mathbf{U}_{ij}$ depends on $\mathbf{X}_i$, and derived other features of $\mathbf{U}_{ij}$ like skewness, multimodality, heavy-tails etc. from the scaled errors $\boldsymbol{\epsilon}_{ij}$. As shown in Sarkar, et al. (2014), even in the much simpler univariate set up, in the absence of precise information on $X_{i\ell}$, variations in other features of $U_{ij\ell}$ for varying values of $X_{i\ell}$, if any, are extremely difficult to detect. More importantly, semiparametric methods that make the multiplicative structural assumption $(U_{ij\ell} \mid X_{i\ell}) = s_\ell(X_{i\ell})\epsilon_{ij\ell}$ are highly robust to departures from this assumption and significantly outperform possible non-parametric alternatives that allow all order moments of $U_{ij\ell}$ to vary flexibly with $X_{i\ell}$, not just the conditional variance, even in scenarios where the true data generating process closely conforms to these nonparametric alternatives.

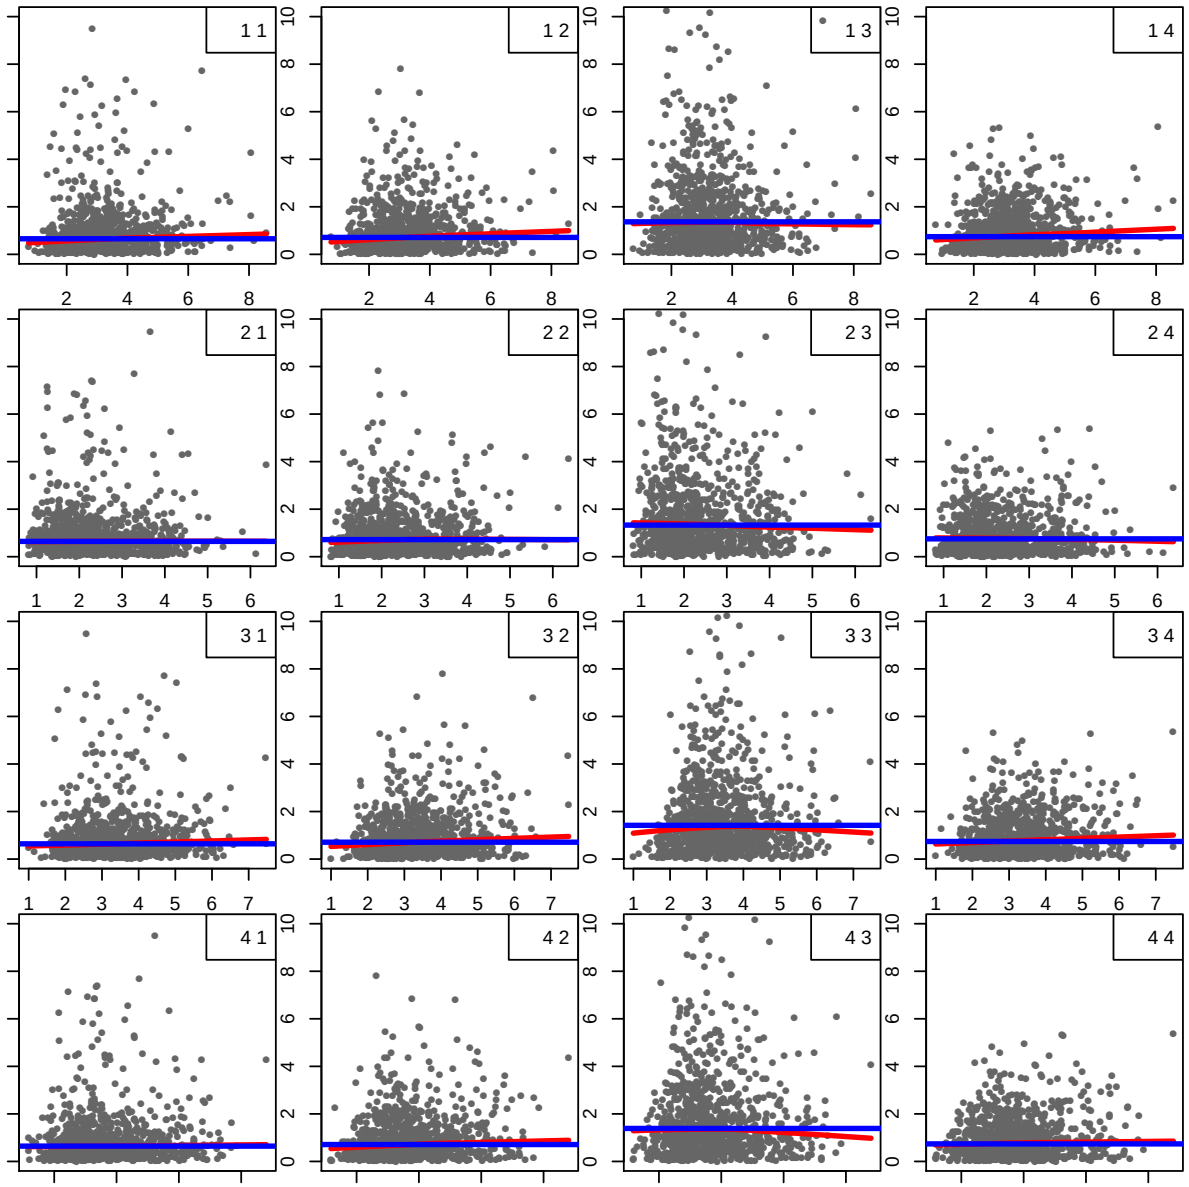

Figure S.1: Panel (ℓ, m) shows the plot of estimates $\widehat{X}_{i\ell}$ of $X_{i\ell}$ vs subject specific variances $\widehat{S}^2_{\epsilon,im}$ of scaled measurement error residuals $\widehat{\epsilon}_{ijm}$, produced by univariate deconvolution methods. See Section S.3 of the Supplementary Materials for additional details. The darker horizontal lines in each panel represent the upper 10% trimmed mean of the subject specific variances $\widehat{S}^2_{\epsilon,i\ell}$. The lighter solid lines in each panel represent nonparametric lowess fits.

S.4 Proofs of Theoretical Results

S.4.1 Proof of Lemma 1

Proof of part 1 of Lemma 1 follows mostly by modifications of the results of Norets and Pelenis (2012). We present here only the proof of part 2 that requires additional modifications along the lines of Pelenis (2014) to accommodate the mean zero restriction on the density of the measurement errors. The first step is to construct finite mixture models of the form

$$f_m(\mathbf{z} \mid \boldsymbol{\theta}_m) = \sum_{k=1}^{m+2} \pi_{m,k} \ \mathrm{MVN}_p(\mathbf{z} \mid \boldsymbol{\mu}_{m,k}, \boldsymbol{\Sigma}_{m,k}) \quad \text{with} \quad \sum_{k=1}^{m+2} \pi_{m,k}\boldsymbol{\mu}_{m,k} = \mathbf{0}$$

that can approximate any given density f_0 that has mean zero and satisfies Conditions 1 with any desired level of accuracy. The continuity of $f_m(\cdot \mid \boldsymbol{\theta})$ implies that the KL distance between f_0 and f_m remains small on sufficiently small open neighborhoods around $\boldsymbol{\theta}_m$. Both the MIW and the MLFA priors assign positive probability to open neighborhoods around $\boldsymbol{\theta}_m$. The conclusion of part 2 of Lemma 1 follows since the prior probability of having $(m+2)$ mixture components is also positive for all $m \in \mathbb{N}$.

Lemma 5. *For any $f_0 \in \widetilde{\mathcal{F}}_{\boldsymbol{\epsilon}}$ and $\eta > 0$, there exists $\boldsymbol{\theta}_m$ such that $d_{KL}\{f_0(\cdot), f_m(\cdot|\boldsymbol{\theta}_m)\} < \eta$.*

Proof. Let $\{A_{m,k}\}_{k=1}^m$ be adjacent cubes with side length h_m, and $A_{m,0} = \mathbb{R}^p - \cup_{k=1}^m A_{m,k}$ such that $h_m \downarrow 0$ but $\cup_{k=1}^m A_{m,k} \uparrow \mathbb{R}^p$ as $m \to \infty$. So $\{A_{m,k}\}_{k=1}^m$ becomes finer but $\cup_{k=1}^m A_{m,k}$ covers more of $\mathbb{R}^p$. Additionally, let the partition be constructed in such a way that for all m sufficiently large, if $\boldsymbol{\epsilon} \in A_{m,0}$, then $C_r(\boldsymbol{\epsilon}) \cap A_{m,0}$ contains a hypercube $C_0(\boldsymbol{\epsilon})$ with side length $r/2$ and a vertex at $\boldsymbol{\epsilon}$; and if $\boldsymbol{\epsilon} \notin A_{m,0}$, then $C_r(\boldsymbol{\epsilon}) \cap (\mathbb{R}^p - A_{m,0})$ contains a hypercube $C_1(\boldsymbol{\epsilon})$ with side length $r/2$ and a vertex at $\boldsymbol{\epsilon}$. Consider the model

$$f_m(\mathbf{z}) \ = \ f_m(\mathbf{z} \mid \boldsymbol{\theta}_m) = \sum_{k=1}^{m+2} \pi_{m,k} \ \mathrm{MVN}_p(\mathbf{z} \mid \boldsymbol{\mu}_{m,k}, \boldsymbol{\Sigma}_{m,k}).$$

Set $\pi_{m,k} = \int_{A_{m,k}} f_0(\mathbf{z})d\mathbf{z}$ for $k = 1, 2, \ldots, m$ and $\pi_{m,k} = P_{f_0}(A_{m,0})/2 = \int_{A_{m,k}} f_0(\mathbf{z})d\mathbf{z}/2$ for $k = (m+1), (m+2)$. Then $\sum_{k=1}^{m+2} \pi_{m,k} = \int_{\mathbb{R}^p} f_0(\mathbf{z})d\mathbf{z} = 1$. Define $g(\mathbf{d}) = \sum_{k=1}^m \pi_{m,k}(\mathbf{c}_{m,k} + \mathbf{d}) + \int_{A_{m,0}} \mathbf{z} f_0(\mathbf{z})d\mathbf{z}$, where $\mathbf{c}_{m,k}$ is the center of $A_{m,k}$ for $k = 1, 2, \ldots, m$.

$$\begin{aligned} g(h_m\mathbf{1}_p/2) \ &= \ \sum_{k=1}^m \pi_{m,k}(\mathbf{c}_{m,k} + h_m\mathbf{1}_p/2) + \int_{A_{m,0}} \mathbf{z} f_0(\mathbf{z})d\mathbf{z} \\ &= \ \sum_{k=1}^m \int_{A_{m,k}} (\mathbf{c}_{m,k} + h_m\mathbf{1}_p/2) f_0(\mathbf{z})d\mathbf{z} + \int_{A_{m,0}} \mathbf{z} f_0(\mathbf{z})d\mathbf{z} \\ &\geq \ \sum_{k=1}^m \int_{A_{m,k}} \mathbf{z} f_0(\mathbf{z})d\mathbf{z} + \int_{A_{m,0}} \mathbf{z} f_0(\mathbf{z})d\mathbf{z} = \int_{\mathbb{R}^p} \mathbf{z} f_0(\mathbf{z})d\mathbf{z} = \mathbf{0}. \end{aligned}$$

Similarly $g(-h_m\mathbf{1}_p/2) \leq 0$. Since $g(\cdot)$ is continuous, there exists $\mathbf{d}_m \in [-h_m/2, h_{m/2}]^p$ such that $g(\mathbf{d}_m) = \mathbf{0}$. Set $\boldsymbol{\mu}_{m,k} = (\mathbf{c}_{m,k} + \mathbf{d}_m)$ for $k = 1, 2, \ldots, m$. Also set $\boldsymbol{\mu}_{m,m+1} =$

$2\int_{A_{m,0}} \mathbf{z} f_0(\mathbf{z})d\mathbf{z} / \int_{A_{m,0}} f_0(\mathbf{z})d\mathbf{z}$ and $\boldsymbol{\mu}_{m,m+2} = \mathbf{0}$ when $\int_{A_{m,0}} f_0(\mathbf{z})d\mathbf{z} > 0$, and $\boldsymbol{\mu}_{m,0} = \mathbf{0}$ otherwise. Then $\sum_{k=1}^{m+2} \pi_{m,k}\boldsymbol{\mu}_{m,k} = g(\mathbf{d}_m) = \mathbf{0}$. Also set $\boldsymbol{\Sigma}_{m,k} = \sigma_m^2 \mathrm{I}_p$ for $k = 1, 2, \ldots, m$ with $\sigma_m \to 0$, and $\Sigma_{m,m+1} = \Sigma_{m,m+2} = \sigma_0^2 \mathrm{I}_p$.

Consider a sequence $\{\delta_m\}_{m=1}^{\infty}$ satisfying $\delta_m > 6p^{1/2}h_m$ and $\delta_m \to 0$. Fix $\boldsymbol{\epsilon} \in \mathbb{R}^p$. Define $C_{\delta_m}(\boldsymbol{\epsilon}) = [\boldsymbol{\epsilon} - \delta_m \mathbf{1}_p/2, \boldsymbol{\epsilon} + \delta_m \mathbf{1}_p/2]$. For m sufficiently large $C_{\delta_m}(\boldsymbol{\epsilon}) \subseteq \cup_{k=1}^{m} A_{m,k}$, $C_{\delta_m}(\boldsymbol{\epsilon}) \cap A_{m,0} = \phi$ and the set $\{k : 1 \le k \le m, A_{m,k} \subset C_{\delta_m}(\boldsymbol{\epsilon})\}$ is non-empty. For $k = 1, \ldots, m$, when $A_{m,k} \subset C_{\delta_m}(\boldsymbol{\epsilon})$, $\pi_{m,k} \ge \inf_{\mathbf{z} \in C_{\delta_m}(\boldsymbol{\epsilon})} f_0(\mathbf{z}) h_m^p$. Therefore,

$$\begin{aligned} f_m(\boldsymbol{\epsilon}) &\ge \sum_{\{k:1\le k\le m, A_{m,k}\subset C_{\delta_m}(\boldsymbol{\epsilon})\}} \pi_{m,k}\ \mathrm{MVN}_p(\boldsymbol{\epsilon} \mid \boldsymbol{\mu}_{m,k}, \sigma_m^2 \mathrm{I}_p) \\ &\ge \inf_{z \in C_{\delta_m}(\boldsymbol{\epsilon})} f_0(\mathbf{z}) \sum_{\{k: A_{m,k}\subset C_{\delta_m}(\boldsymbol{\epsilon})\}} h_m^p\ \mathrm{MVN}_p(\boldsymbol{\epsilon} \mid \mathbf{c}_{m,k} + \mathbf{d}_m, \sigma_m^2 \mathrm{I}_p) \\ &\ge \inf_{z \in C_{\delta_m}(\boldsymbol{\epsilon})} f_0(\mathbf{z}) \quad \left\{1 - \frac{6p^{3/2}h_m\delta_m^{p-1}}{(2\pi)^{p/2}\sigma_m^p} - \frac{8p\sigma_m}{(2\pi)^{1/2}\delta_m}\right\}, \end{aligned}$$

where the last step follows from Lemma 1 and Lemma 2 of Norets and Pelenis (2012). Let h_m, δ_m, σ_m further satisfy $h_m/\sigma_m^p \to 0, \sigma_m/\delta_m \to 0$. Then for any $\eta > 0$ there exists an M_1 large enough such that for all $m > M_1$

$$f_m(\boldsymbol{\epsilon}) \ge \inf_{\mathbf{z} \in C_{\delta_m}(\boldsymbol{\epsilon})} f_0(\mathbf{z}) \cdot (1 - \eta).$$

Without loss of generality, we may assume $f_0(\boldsymbol{\epsilon}) > 0$. Since $f_0(\cdot)$ is continuous and $\delta_m \to 0$, there also exists an M_2 such that for all $m > M_2$ we have $\inf_{\mathbf{z} \in C_{\delta_m}(\boldsymbol{\epsilon})} f_0(\mathbf{z}) > 0$ and

$$\frac{f_0(\boldsymbol{\epsilon})}{\inf_{\mathbf{z} \in C_{\delta_m}(\boldsymbol{\epsilon})} f_0(\mathbf{z})} \le (1 + \eta).$$

Therefore, for all $m > \max\{M_1, M_2\}$, we have

$$1 \le \max\left\{1, \frac{f_0(\boldsymbol{\epsilon})}{f_m(\boldsymbol{\epsilon})}\right\} \le \max\left\{1, \frac{f_0(\boldsymbol{\epsilon})}{\inf_{z \in C_{\delta_m}(\boldsymbol{\epsilon})} f_0(z) \cdot (1 - \eta)}\right\} \le \frac{(1 + \eta)}{(1 - \eta)}.$$

Thus, $\log \max\{1, f_0(\boldsymbol{\epsilon})/f_m(\boldsymbol{\epsilon})\} \to 0$ as $m \to \infty$. Pointwise convergence is thus established. Next, we will find an integrable upper bound for $\log \max\{1, f_0(\boldsymbol{\epsilon})/f_m(\boldsymbol{\epsilon})\}$.

For point wise convergence we can assume $\boldsymbol{\epsilon} \notin A_{m,0}$ for sufficiently large m. But to find integrable upper bound, we have to consider both the cases $\boldsymbol{\epsilon} \in A_{m,0}$ and $\boldsymbol{\epsilon} \notin A_{m,0}$. When $\boldsymbol{\epsilon} \in A_{m,0}$, we have $\mathrm{P}_{f_0}(A_{m,0}) = \int_{A_{m,0}} f_0(\mathbf{z})d\mathbf{z} \ge \int_{A_{m,0}\cap C_r(\boldsymbol{\epsilon})} f_0(\mathbf{z})d\mathbf{z} \ge \lambda\{A_{m,0} \cap C_r(\boldsymbol{\epsilon})\} \inf_{\mathbf{z} \in A_{m,0}\cap C_r(\boldsymbol{\epsilon})} f_0(\mathbf{z}) \ge (r/2)^p \inf_{\mathbf{z} \in C_r(\boldsymbol{\epsilon})} f_0(\mathbf{z})$, since $\lambda\{A_{m,0} \cap C_r(\boldsymbol{\epsilon})\} \ge \lambda\{C_0(\boldsymbol{\epsilon})\} \ge (r/2)^p$. Using part 4 of Conditions 1 and Lemma 1 and Lemma 2 of Norets and Pelenis (2012) again, if $\boldsymbol{\epsilon} \notin A_{m,0}$, for m sufficiently large

$$\sum_{\{k: A_{m,k}\subset C_r(\boldsymbol{\epsilon})\}} h_m^p\ \mathrm{MVN}_p(\boldsymbol{\epsilon} \mid \boldsymbol{\mu}_{m,k}, \sigma_m^2 \mathrm{I}_p) \ge \sum_{\{k: A_{m,k}\subset C_1(\boldsymbol{\epsilon})\}} h_m^p\ \mathrm{MVN}_p(\boldsymbol{\epsilon} \mid \boldsymbol{\mu}_{m,k}, \sigma_m^2 \mathrm{I}_p)$$

$$\geq \int_{C_1(\boldsymbol{\epsilon})} \text{MVN}_p(\mathbf{z} \mid \boldsymbol{\epsilon}, \sigma_m^2 \text{I}_p) d\mathbf{z} - \frac{3p^{3/2}(r/2)^{p-1} h_m}{(2\pi)^{p/2}\sigma_m^p}$$
$$\geq \left\{ \frac{1}{2^p} - \frac{8p\sigma_m}{2^p(2\pi)^{1/2} r} - \frac{3p^{3/2} h_m r^{p-1}}{2^{p-1}(2\pi)^{p/2}\sigma_m^p} \right\} \geq \frac{1}{2^{p+1}},$$

This implies

$$\begin{aligned}
f_m(\boldsymbol{\epsilon}) &= \sum_{k=1}^{m} P_{f_0}(A_{m,k})\ \text{MVN}_p(\boldsymbol{\epsilon} \mid \boldsymbol{\mu}_{m,k}, \sigma_m^2 \text{I}_p) + \sum_{k=m+1}^{m+2} (1/2) P_{f_0}(A_{m,0})\ \text{MVN}_p(\boldsymbol{\epsilon} \mid \boldsymbol{\mu}_{m,k}, \sigma_0^2 \text{I}_p) \\
&\geq \sum_{k=1}^{m} P_{f_0}(A_{m,k})\ \text{MVN}_p(\boldsymbol{\epsilon} \mid \boldsymbol{\mu}_{m,k}, \sigma_m^2 \text{I}_p) + (1/2) P_{f_0}(A_{m,0})\ \text{MVN}_p(\boldsymbol{\epsilon} \mid \mathbf{0}, \sigma_0^2 \text{I}_p) \\
&\geq \{1 - 1(\boldsymbol{\epsilon} \in A_{m,0})\} \inf_{\mathbf{z} \in C_r(\boldsymbol{\epsilon})} f_0(\mathbf{z}) \sum_{\{k: A_{m,k} \subset C_r(\boldsymbol{\epsilon})\}} \lambda(A_{m,k})\ \text{MVN}_p(\boldsymbol{\epsilon} \mid \boldsymbol{\mu}_{m,k}, \sigma_m^2 \text{I}_p) \\
&\quad + 1(\boldsymbol{\epsilon} \in A_{m,0}) (1/2) P_{f_0}(A_{m,0})\ \text{MVN}_p(\boldsymbol{\epsilon} \mid \mathbf{0}, \sigma_0^2 \text{I}_p) \\
&\geq (1/2)\{1 - 1(\boldsymbol{\epsilon} \in A_{m,0})\} \inf_{\mathbf{z} \in C_r(\boldsymbol{\epsilon})} f_0(\mathbf{z}) \\
&\quad + 1(\boldsymbol{\epsilon} \in A_{m,0})\ (1/2)(r/2)^p\ \text{MVN}_p(\boldsymbol{\epsilon} \mid \mathbf{0}, \sigma_0^2 \text{I}_p) \inf_{\mathbf{z} \in C_r(\boldsymbol{\epsilon})} f_0(\mathbf{z}) \\
&\geq (1/2)(r/2)^p\ \text{MVN}_p(\boldsymbol{\epsilon} \mid \mathbf{0}, \sigma_0^2 \text{I}_p) \inf_{\mathbf{z} \in C_r(\boldsymbol{\epsilon})} f_0(\mathbf{z}).
\end{aligned}$$

The last step followed by choosing σ_0^2 large enough so that $(r/2)^p \sup_{\boldsymbol{\epsilon} \in \mathbb{R}^p} \text{MVN}_p(\boldsymbol{\epsilon} \mid \mathbf{0}, \sigma_0^2 \text{I}_p) < (r/2)^p\ \sigma_0^{-p} < 2^{-(p+1)} < 1$. Therefore,

$$\begin{aligned}
&\log \max\left\{1, \frac{f_0(\boldsymbol{\epsilon})}{f_m(\boldsymbol{\epsilon})}\right\} \leq \log \max\left\{1, \frac{f_0(\boldsymbol{\epsilon})}{(1/2)(r/2)^p\ \text{MVN}_p(\boldsymbol{\epsilon} \mid \mathbf{0}, \sigma_0^2 \text{I}_p)\ \inf_{\mathbf{z} \in C_r(\boldsymbol{\epsilon})} f_0(\mathbf{z})}\right\} \\
&\leq \log \frac{1}{(1/2)(r/2)^p\ \text{MVN}_p(\boldsymbol{\epsilon} \mid \mathbf{0}, \sigma_0^2 \text{I}_p)} \max\left\{(1/2)(r/2)^p\ \text{MVN}_p(\boldsymbol{\epsilon} \mid \mathbf{0}, \sigma_0^2 \text{I}_p), \frac{f_0(\boldsymbol{\epsilon})}{\inf_{\mathbf{z} \in C_r(\boldsymbol{\epsilon})} f_0(\mathbf{z})}\right\} \\
&\leq -\log\left\{(1/2)(r/2)^p\ \text{MVN}_p(\boldsymbol{\epsilon} \mid \mathbf{0}, \sigma_0^2 \text{I}_p)\right\} + \log\left\{\frac{f_0(\boldsymbol{\epsilon})}{\inf_{\mathbf{z} \in C_r(\boldsymbol{\epsilon})} f_0(\mathbf{z})}\right\}.
\end{aligned}$$

The first and the second terms are integrable by part 2 and part 3 of Conditions 1, respectively. Since $\int f_0(\boldsymbol{\epsilon}) \log\{f_0\boldsymbol{\epsilon}/f_m(\boldsymbol{\epsilon})\} d\boldsymbol{\epsilon} \leq \int f_0(\boldsymbol{\epsilon}) \log \max\{1, f_0\boldsymbol{\epsilon}/f_m(\boldsymbol{\epsilon})\} d\boldsymbol{\epsilon}$, the proof of Lemma 5 is completed applying dominated convergence theorem (DCT).

Let $\eta > 0$ be given. According to Lemma 5, there exists $\boldsymbol{\theta}_m^\star = (\boldsymbol{\pi}_{1:(m+2)}^\star, \boldsymbol{\mu}_{1:(m+2)}^\star, \boldsymbol{\Sigma}_{1:(m+2)}^\star)$ with $\boldsymbol{\Sigma}_k^\star = \sigma_m^{2\star} \text{I}_p$ for $k = 1, \ldots, m$ and $\boldsymbol{\Sigma}_k^\star = \sigma_0^{2\star} \text{I}_p$ for $k = (m+1), (m+2)$ such that $d_{KL}\{f_0(\cdot), f_m(\cdot \mid \boldsymbol{\theta}_m^\star)\} < \eta/2$. We have, for any $\boldsymbol{\theta}_m$,

$$\int f_0(\boldsymbol{\epsilon})\ \log\left\{\frac{f_0(\boldsymbol{\epsilon})}{f_m(\boldsymbol{\epsilon} \mid \boldsymbol{\theta}_m)}\right\} d\boldsymbol{\epsilon} = \int f_0(\boldsymbol{\epsilon})\ \log\left\{\frac{f_0(\boldsymbol{\epsilon})}{f_m(\boldsymbol{\epsilon} \mid \boldsymbol{\theta}_m^\star)}\right\} d\boldsymbol{\epsilon} + \int f_0(\boldsymbol{\epsilon})\ \log\left\{\frac{f_m(\boldsymbol{\epsilon} \mid \boldsymbol{\theta}_m^\star)}{f_m(\boldsymbol{\epsilon} \mid \boldsymbol{\theta}_m)}\right\} d\boldsymbol{\epsilon}.$$

Let the second term in the above expression be denoted by $g(\boldsymbol{\theta}_m)$. The priors puts positive mass on arbitrarily small open neighborhoods around $\boldsymbol{\theta}_m^\star$. The result will follow if

there exists an open neighborhood $\mathcal{N}(\boldsymbol{\theta}_m^\star)$ around $\boldsymbol{\theta}_m^\star$ such that $\sup_{\boldsymbol{\theta}_m \in \mathcal{N}(\boldsymbol{\theta}_m^\star)} g(\boldsymbol{\theta}_m) < \eta/2$. Since $g(\boldsymbol{\theta}_m^\star) = 0$, it suffices to show that the function $g(\boldsymbol{\theta}_m)$ is continuous at $\boldsymbol{\theta}_m^\star$. Now $g(\boldsymbol{\theta})$ is continuous at $\boldsymbol{\theta}_m^\star$ if for every sequence $\{\boldsymbol{\theta}_{m,n}\}_{n=1}^\infty$ with $\boldsymbol{\theta}_{m,n} \to \boldsymbol{\theta}_m^\star$, we have $g(\boldsymbol{\theta}_{m,n}) \to g(\boldsymbol{\theta}_m^\star)$. For all $\boldsymbol{\epsilon} \in \mathbb{R}^p$, we have $\log\{f_m(\boldsymbol{\epsilon} \mid \boldsymbol{\theta}_{m,n}^\star)/f_m(\boldsymbol{\epsilon} \mid \boldsymbol{\theta}_m)\} \to 0$ as $\boldsymbol{\theta}_{m,n} \to \boldsymbol{\theta}_m^\star$. Continuity of $g(\boldsymbol{\theta}_m)$ at $\boldsymbol{\theta}_m^\star$ will follow from DCT if we can show that $|f_m(\boldsymbol{\epsilon} \mid \boldsymbol{\theta}_m^\star)/f_m(\boldsymbol{\epsilon} \mid \boldsymbol{\theta}_{m,n})|$ has an integrable with respect to f_0 upper bound.

Since $\boldsymbol{\theta}_{m,n} \to \boldsymbol{\theta}_m^\star$, for any arbitrarily small open neighborhood $\mathcal{N}(\boldsymbol{\theta}_m^\star)$ around $\boldsymbol{\theta}_m^\star$, we must have $\boldsymbol{\theta}_{m,n} \in \mathcal{N}(\boldsymbol{\theta}_m^\star)$ for all n sufficiently large. Let $\boldsymbol{\theta}_m = (\boldsymbol{\pi}_{1:(m+2)}, \boldsymbol{\mu}_{1:(m+2)}, \boldsymbol{\Sigma}_{1:(m+2)}) \in \mathcal{N}(\boldsymbol{\theta}_m^\star)$. Since the eigenvalues of a real symmetric matrix depend continuously on the matrix, we must have $(\lambda_1(\boldsymbol{\Sigma}_k), \lambda_p(\boldsymbol{\Sigma}_k)) \subset (\underline{\sigma}_m^{2\star}, \overline{\sigma}_m^{2\star})$ for $k = 1, \ldots, m$ and $(\lambda_1(\boldsymbol{\Sigma}_k), \lambda_p(\boldsymbol{\Sigma}_k)) \subset (\underline{\sigma}_0^{2\star}, \overline{\sigma}_0^{2\star})$ for $k = (m+1), (m+2)$, where $\underline{\sigma}_m^{2\star} < \sigma_m^{2\star} < \overline{\sigma}_m^{2\star}$ and $\underline{\sigma}_0^{2\star} < \sigma_0^{2\star} < \overline{\sigma}_0^{2\star}$. Let $\underline{\sigma}^{2\star} = \min\{\underline{\sigma}_m^{2\star}, \underline{\sigma}_0^{2\star}\}$ and $\overline{\sigma}^{2\star} = \max\{\overline{\sigma}_m^{2\star}, \overline{\sigma}_0^{2\star}\}$. Then $(\lambda_1(\boldsymbol{\Sigma}_k), \lambda_p(\boldsymbol{\Sigma}_k)) \subset (\underline{\sigma}^{2\star}, \overline{\sigma}^{2\star})$ for $k = 1, \ldots, (m+2)$. Similarly, for some finite $\mu^\star$, we must have $\boldsymbol{\mu}_{m,k} \in (-\mu^\star \mathbf{1}_p, \mu^\star \mathbf{1}_p) = \mathcal{N}_{\mu^\star}$ for $k = 1, \ldots, (m+2)$. For any real positive definite matrix $\boldsymbol{\Sigma}$, we have $\mathbf{z}^{\mathrm{T}}\boldsymbol{\Sigma}^{-1}\mathbf{z} \leq \lambda_1^{-1}(\boldsymbol{\Sigma})\|\mathbf{z}\|^2$. Therefore, for any $\boldsymbol{\epsilon} \in \mathbb{R}^p$ and for all $k = 1, \ldots, (m+2)$, we must have $(\boldsymbol{\epsilon} - \boldsymbol{\mu}_{m,k})^{\mathrm{T}}\boldsymbol{\Sigma}_{m,k}^{-1}(\boldsymbol{\epsilon} - \boldsymbol{\mu}_{m,k}) \leq \underline{\sigma}^{-2\star}\{1(\boldsymbol{\epsilon} \in \mathcal{N}_{\mu^\star})2^p\mu^{\star p} + 1(\boldsymbol{\epsilon} \notin \mathcal{N}_{\mu^\star})\|\boldsymbol{\epsilon} + \mathrm{sign}(\boldsymbol{\epsilon})\mu^\star\|^2\}$, where $\mathrm{sign}(\boldsymbol{\epsilon}) = \{\mathrm{sign}(\epsilon_1), \ldots, \mathrm{sign}(\epsilon_p)\}^{\mathrm{T}}$. Therefore, for any $\boldsymbol{\theta}_m \in \mathcal{N}(\boldsymbol{\theta}_m^\star)$, we have

$$[1(\boldsymbol{\epsilon} \in \mathcal{N}_{\mu^\star})\mathrm{MVN}_p(2\mu^\star \mathbf{1}_p \mid \mathbf{0}, \underline{\sigma}^{2\star}\mathrm{I}_p) + 1(\boldsymbol{\epsilon} \notin \mathcal{N}_{\mu^\star})\mathrm{MVN}_p\{\boldsymbol{\epsilon} + \mathrm{sign}(\boldsymbol{\epsilon})\mu^\star \mid \mathbf{0}, \underline{\sigma}^{2\star}\mathrm{I}_p\}]/\overline{\sigma}^\star \leq f_m(\boldsymbol{\epsilon} \mid \boldsymbol{\theta}_m) \leq 1/\underline{\sigma}^\star.$$

The upper bound is a constant and the logarithm of the lower bound is integrable since, by part 2 of Conditions 1, the second order moments of $\boldsymbol{\epsilon}$ exist. An f_0 integrable upper bound for the function $\sup_{\boldsymbol{\theta}_m \in \mathcal{N}(\boldsymbol{\theta}_m^\star)} |f_m(\boldsymbol{\epsilon} \mid \boldsymbol{\theta}_m)|$ thus exists. Finally, DCT applies because

$$\begin{aligned}\int f_0(\boldsymbol{\epsilon}) \left|\log\left\{\frac{f_m(\boldsymbol{\epsilon} \mid \boldsymbol{\theta}_m^\star)}{f_m(\boldsymbol{\epsilon} \mid \boldsymbol{\theta}_{m,n})}\right\}\right| d\boldsymbol{\epsilon} &\leq \sup_{\boldsymbol{\theta}_m \in \mathcal{N}(\boldsymbol{\theta}_m^\star)} \int f_0(\boldsymbol{\epsilon}) \left|\log\left\{\frac{f_m(\boldsymbol{\epsilon} \mid \boldsymbol{\theta}_m^\star)}{f_m(\boldsymbol{\epsilon} \mid \boldsymbol{\theta}_m)}\right\}\right| d\boldsymbol{\epsilon} \\ &\leq 2 \sup_{\boldsymbol{\theta}_m \in \mathcal{N}(\boldsymbol{\theta}_m^\star)} \int f_0(\boldsymbol{\epsilon}) \; |f_m(\boldsymbol{\epsilon} \mid \boldsymbol{\theta}_m)| \, d\boldsymbol{\epsilon}.\end{aligned}$$

The conclusion of part 2 of Lemma 1 follows since the prior probability of having $(m+2)$ mixture components is positive for all $m \in \mathbb{N}$.

S.4.2 Proof of Lemma 2

Given q, let Π_q denote a prior on $\mathbb{N}_q = \{q+1, q+2, \ldots\}$ such that $\Pi_q(J) > 0 \; \forall J \in \mathbb{N}_q$. Let $||\cdot||_2$ denote the Euclidean norm. Let $\mathbb{R}^+ = (0, \infty)$. Given $J \sim \Pi_q$, also let $\Pi_{\beta|J}$ be a prior on $\mathbb{R}^{+J}$ such that $\Pi_{\beta|J}\{N_\delta(\boldsymbol{\beta}_0)\} > 0$ for any $\delta > 0$ and any $\boldsymbol{\beta}_0 \in \mathbb{R}^J$, where $N_\delta(\boldsymbol{\beta}_0) = \{\boldsymbol{\beta} : \boldsymbol{\beta} \in \mathbb{R}^{+J}, ||\boldsymbol{\beta} - \boldsymbol{\beta}_0||_2 < \delta\}$. Define $\mathcal{S}_{q,J} = \{v_s : v_s = \mathbf{B}_{q,J}\boldsymbol{\beta} = \sum_{j=1}^J b_{q,j}\beta_j \text{ for some } \boldsymbol{\beta} \in \mathbb{R}^{+J}\}$. Then $\Pi_{\mathbf{V}} = \Pi_q \times \Pi_{\beta|J}$ is the induced prior on $\mathcal{S}_q = \cup_{J=q+1}^\infty \mathcal{S}_{q,J}$.

Define $\psi(v_0, h) = \sup_{X, X' \in [A,B], |X - X'| \leq h} |v_0(X) - v_0(X')|$. Let $\lfloor \alpha \rfloor = \min\{n : n \in \mathbb{N}, n \geq$

$\alpha\}$. For any X, (i) $b_{q,j}(X) \geq 0\ \forall j$, (ii) $\sum_{j=1}^{J} b_{q,j}(X) = 1$, (iii) $b_{q,j}$ is positive only inside the interval $[t_j, t_{j+q+1}]$, and (iv) for $j \in \{(q+1), (q+2), \ldots, (q+K)\}$, for any $X \in (t_j, t_{j+1})$, only $(q+1)$ B-splines $b_{q,j-q}(X), b_{q,j-q+1}(X), \ldots, b_{q,j}(X)$ are positive. Using these local support properties of B-splines, the results on page 147 of de Boor (2000) can be modified to show that, for any $v_0 \in \mathcal{C}_+[A, B]$,

$$\inf_{v_s \in \mathcal{S}_{q,J}} ||v_0 - v_s||_\infty \leq \lfloor (q+1)/2 \rfloor\ \psi(v_0, \Delta_{\max}) \to 0 \ \text{ as } \Delta_{\max} \to 0.$$

Also, if $q \geq (\alpha - 1)$, we can modify the results on page 149 of de Boor (2000) to show that, for any $v \in \mathcal{C}_+^\alpha[A, B]$,

$$\inf_{v_s \in \mathcal{S}_{q,J}} ||v_0 - v_s||_\infty \ \ \leq \ \ c(q)c(q-1)\ldots c(q-\alpha_0+1)\ ||v_0^{(\alpha_0)}||_\infty\ \Delta_{\max}^{\alpha_0},$$

where $c(q) = \lfloor (q+1)/2 \rfloor$. For any two functions g_1 and g_2, $\sup|g_1 g_2| \leq \sup|g_1| \sup|g_2|$. Taking $g_1(X, X') = \{v_0^{(\alpha_0)}(X) - v_0^{(\alpha_0)}(X')\}/(X - X')^{(\alpha - \alpha_0)}$ and $g_2(X, X') = (X - X')^{(\alpha - \alpha_0)}$, we have $||v_0^{(\alpha_0)}||_\infty \leq ||v_0||_\alpha (B - A)^{(\alpha - \alpha_0)}$. Therefore,

$$\inf_{v_s \in \mathcal{S}_{q,J}} ||v_0 - v_s||_\infty \ \ \leq \ \ c(q, \alpha_0)\ (B - A)^{(\alpha - \alpha_0)}\ ||v_0||_\alpha\ \Delta_{\max}^{\alpha_0}.$$

Furthermore, when the knot points $\{t_{q+1+j}\}_{j=0}^{K}$ are equidistant

$$\inf_{v_s \in \mathcal{S}_{q,J}} ||v_0 - v_s||_\infty \leq c(q, \alpha_0) ||v_0^{(\alpha)}||_\infty \frac{(B-A)^\alpha}{K^{\alpha_0}} \leq c(q, \alpha) ||v_0||_\alpha K^{-\alpha}.$$

Given any $v_0 \in C_+[A, B]$(or $C_+^\alpha[A, B]$) and $\delta > 0$, find $J \in \mathbb{N}_q$ and $\boldsymbol{\beta}_0 \in \mathbb{R}^{+J}$ such that $||v_0 - \mathbf{B}_{q,J}\boldsymbol{\beta}_0||_\infty = \inf_{v_s \in \mathcal{S}_{q,J}} ||v_0 - v_s||_\infty < \delta/2$. Next consider a neighborhood $N_\eta(\boldsymbol{\beta}_0)$ such that for any $\boldsymbol{\beta} \in N_\eta(\boldsymbol{\beta}_0)$, we have $||\mathbf{B}_{q,J}\boldsymbol{\beta} - \mathbf{B}_{q,J}\boldsymbol{\beta}_0||_\infty < \delta/2$. Then for any $\boldsymbol{\beta} \in N_\eta(\boldsymbol{\beta}_0)$, we have $||\mathbf{B}_{q,J}\boldsymbol{\beta} - v_0||_\infty \leq ||\mathbf{B}_{q,J}\boldsymbol{\beta} - \mathbf{B}_{q,J}\boldsymbol{\beta}_0||_\infty + ||\mathbf{B}_{q,J}\boldsymbol{\beta}_0 - v_0||_\infty < \delta$. Also $\Pi_{\mathbf{V}}(||v - v_0||_\infty < \delta) \geq \Pi_q(J)\ \Pi_{\beta|J}\{N_\eta(\boldsymbol{\beta}_0)\} > 0$. Proof of Lemma 2 then follows as a special case taking $\boldsymbol{\beta} = \exp(\boldsymbol{\xi})$ and taking Π_q and $\Pi_{\beta|J}$ to be the priors on J and $\boldsymbol{\beta}$ induced by $p_0(K)$ and $p_0(\boldsymbol{\xi}|K, \sigma_\xi^2)$, respectively.

S.4.3 Proof of Lemma 3

We first prove some additional lemmas to used in the proof of Lemma 3.

Lemma 6. $\Pi_{\mathbf{V}}(||v - v_0||_\infty < \delta) > 0\ \forall \delta > 0$ *implies that* $\Pi_{\mathbf{V}}(||g \circ v - g \circ v_0||_\infty < \delta) > 0\ \forall \delta > 0$ *for every continuous function* $g : \mathbb{R} \to \mathbb{R}$.

Proof. Let $v : [A, B] \to [C_1, D_1]$ and $v_0 : [A, B] \to [C_2, D_2]$. Then $(v - v_0) : [A, B] \to [C_1 - D_2, D_1 - C_2] = [C, D]$, say. Then $g : [C, D] \to \mathbb{R}$ is a uniformly continuous function. Therefore, given any $\delta > 0$, there exists a $\eta > 0$ such that $|g(Z_1) - g(Z_2)| < \delta$ whenever $|Z_1 - Z_2| < \eta$. Now let $||v - v_0||_\infty = \sup_{X \in [A,B]} |v(X) - v_0(X)| < \eta$. This implies, for all

$X \in [A, B]$, $|v(X) - v_0(X)| < \eta$. Therefore, for all $X \in [A, B]$, $|g\{v(X)\} - g\{v_0(X)\}| < \delta$, and hence $||g \circ v - g \circ v_0||_\infty \leq \delta$. Hence the proof.

Corollary 1. *In particular, taking $g(Z) = Z^{1/2}$ $\forall Z > 0$ and $g(Z) = 0$ otherwise, we have $\Pi_{\mathbf{V}}(||v^{1/2} - v_0^{1/2}||_\infty < \delta) = \Pi_{\mathbf{V}}(||s - s_0||_\infty < \delta) > 0$ $\forall \delta > 0$ for all $v_0 \in \mathcal{C}_+[A, B]$ (or $\mathcal{C}_+^\alpha[A, B]$).*

Let $P_{\boldsymbol{\epsilon},K}\{(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \mid \boldsymbol{\pi}_{1:K}, \boldsymbol{\mu}_{1:K}, \boldsymbol{\Sigma}_{1:K}\} = \sum_{k=1}^K \pi_k \delta_{(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, where $\delta_{\boldsymbol{\theta}}$ denotes a point mass at $\boldsymbol{\theta}$. Keeping the hyper-parameters implicit, $P_0(\boldsymbol{\pi}_{1:K}, \boldsymbol{\mu}_{1:K}, \boldsymbol{\Sigma}_{1:K}) = P_{0\pi}(\boldsymbol{\pi}_{1:K}) P_{0\mu}(\boldsymbol{\mu}_{1:K} \mid \boldsymbol{\pi}_{1:K}) P_{0\Sigma}(\boldsymbol{\Sigma}_{1:K})$. Denoting $P_{\boldsymbol{\epsilon},K}\{(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \mid \boldsymbol{\pi}_{1:K}, \boldsymbol{\mu}_{1:K}, \boldsymbol{\Sigma}_{1:K}\}$ simply by $P_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Let c be a generic for constants that are not of direct interest. For any square matrix $\mathbf{A}$ of order p, let $\lambda_1(\mathbf{A}) \leq \cdots \leq \lambda_p(\mathbf{A})$ denote the ordered eigenvalues of $\mathbf{A}$. The following lemma proves some properties of $P_{\boldsymbol{\epsilon},K}$ and $f_{\boldsymbol{\epsilon}}$.

Lemma 7. *1. $\int \|\boldsymbol{\mu}\|_2^2 \, dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) < \infty$ a.s. 2. $\int \lambda_1^{-1}(\boldsymbol{\Sigma}) dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) < \infty$ a.s. 3. $\int |\boldsymbol{\Sigma}|^{-1/2} \, dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) < \infty$ a.s.*

Proof. 1. The prior $P_{0\mu}(\boldsymbol{\mu}_{1:K} \mid \boldsymbol{\pi}_{1:K})$ is of the form (15), that is, $P_{0\mu}(\boldsymbol{\mu}_{1:K} \mid \boldsymbol{\pi}_{1:K}) = \text{MVN}_{Kp}(\mathbf{0}, \boldsymbol{\Sigma}^0 - \boldsymbol{\Sigma}^0_{1,R} \boldsymbol{\Sigma}^{-1}_{R,R} \boldsymbol{\Sigma}^0_{R,1})$, where $\boldsymbol{\Sigma}^0$ is a $Kp \times Kp$ block-diagonal matrix independent of $\boldsymbol{\pi}_{1:K}$, all k principal blocks of order $p \times p$ being $\boldsymbol{\Sigma}_0$. The matrix $\boldsymbol{\Sigma}^0_{1,R} \boldsymbol{\Sigma}^{-1}_{R,R} \boldsymbol{\Sigma}^0_{R,1}$ depends on $\boldsymbol{\pi}_{1:K}$ and is nonnegative definite so that its diagonal elements are all nonnegative. Let $\boldsymbol{\Sigma}_0 = ((\sigma_{0,ij}))$ and $\boldsymbol{\Sigma}^0_{1,R} \boldsymbol{\Sigma}^{-1}_{R,R} \boldsymbol{\Sigma}^0_{R,1} = ((\sigma_{R,ij}))$. Then, $\int \|\boldsymbol{\mu}_k\|_2^2 \, dP_{0\mu}(\boldsymbol{\mu}_{1:K} \mid \boldsymbol{\pi}_{1:K}) = \left\{\sum_{j=1}^p \sigma_{0,jj} - \sum_{j=(k-1)p+1}^{kp} \sigma_{R,jj}\right\} \leq \sum_{j=1}^p \sigma_{0,jj} = \text{trace}(\boldsymbol{\Sigma}_0)$. Therefore,

$$\int \int \|\boldsymbol{\mu}\|_2^2 \, dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) dP_0(\boldsymbol{\pi}_{1:K}, \boldsymbol{\mu}_{1:K}, \boldsymbol{\Sigma}_{1:K}) = \sum_{k=1}^K \int \pi_k \|\boldsymbol{\mu}_k\|_2^2 \, dP_{0\mu}(\boldsymbol{\mu}_{1:K} \mid \boldsymbol{\pi}_{1:K}) dP_{0\pi}(\boldsymbol{\pi}_{1:K})$$
$$\leq \text{trace}(\boldsymbol{\Sigma}_0) < \infty.$$

2. We have $\int \int \lambda_1^{-1}(\boldsymbol{\Sigma}) dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) dP_0(\boldsymbol{\pi}_{1:K}, \boldsymbol{\mu}_{1:K}, \boldsymbol{\Sigma}_{1:K}) = \int \lambda_1^{-1}(\boldsymbol{\Sigma}) dP_{0\Sigma}(\boldsymbol{\Sigma})$.

When $\boldsymbol{\Sigma} \sim \text{IW}_p(\nu_0, \boldsymbol{\Psi}_0)$, we have $\boldsymbol{\Psi}_0^{-1/2} \boldsymbol{\Sigma}^{-1} \boldsymbol{\Psi}_0^{-1/2} \sim \text{W}_p(\nu_0, \text{I})$ and $\text{trace}(\boldsymbol{\Psi}_0^{-1} \boldsymbol{\Sigma}^{-1}) = \text{trace}(\boldsymbol{\Psi}_0^{-1/2} \boldsymbol{\Sigma}^{-1} \boldsymbol{\Psi}_0^{-1/2}) \sim \chi^2_{p\nu_0}$. Here $\text{W}_p(\nu, \boldsymbol{\Psi})$ denotes a Wishart distribution with degrees of freedom ν and mean $\nu\boldsymbol{\Psi}$. For any two positive semidefinite matrices $\mathbf{A}$ and $\mathbf{B}$, we have $\lambda_1(\mathbf{A})\text{trace}(\mathbf{B}) \leq \text{trace}(\mathbf{AB}) \leq \lambda_p(\mathbf{A})\text{trace}(\mathbf{B})$. Therefore, $\lambda_1(\boldsymbol{\Psi}_0^{-1}) E\{\text{trace}(\boldsymbol{\Sigma}^{-1})\} \leq E\{\text{trace}(\boldsymbol{\Psi}_0^{-1} \boldsymbol{\Sigma}^{-1})\} = p\nu_0$. Hence, $\int \lambda_1^{-1}(\boldsymbol{\Sigma}) dP_{0\Sigma}(\boldsymbol{\Sigma}) = E\lambda_p(\boldsymbol{\Sigma}^{-1}) \leq E\{\text{trace}(\boldsymbol{\Sigma}^{-1})\} < \infty$.

When $\boldsymbol{\Sigma} = \boldsymbol{\Omega} + \boldsymbol{\Lambda}\boldsymbol{\Lambda}^{\text{T}}$ with $\boldsymbol{\Omega} = \text{diag}(\sigma_1^2, \ldots, \sigma_p^2)$, we have $\text{trace}(\boldsymbol{\Sigma}^{-1}) = \text{trace}\{\boldsymbol{\Omega}^{-1} - \boldsymbol{\Omega}^{-1}\boldsymbol{\Gamma}(\text{I}_p + \boldsymbol{\Gamma}^{\text{T}}\boldsymbol{\Omega}^{-1}\boldsymbol{\Gamma})^{-1}\boldsymbol{\Gamma}^{\text{T}}\boldsymbol{\Omega}^{-1}\} \leq \text{trace}(\boldsymbol{\Omega}^{-1}) = \sum_{j=1}^p \sigma_j^{-2}$, where $\boldsymbol{\Gamma}$ is a $p \times p$ matrix satisfying $\boldsymbol{\Gamma}\boldsymbol{\Gamma}^{\text{T}} = \boldsymbol{\Lambda}\boldsymbol{\Lambda}^{\text{T}}$. Thus, $\int \lambda_1^{-1}(\boldsymbol{\Sigma}) dP_{0\Sigma}(\boldsymbol{\Sigma}_{1:K}) = E\lambda_p(\boldsymbol{\Sigma}^{-1}) \leq E\{\text{trace}(\boldsymbol{\Sigma}^{-1})\} \leq \sum_{j=1}^p E\sigma_j^{-2} < \infty$ whenever $\sigma_j^2 \sim \text{Inv-Ga}(a, b)$ with $a > 1$.

3. When $\boldsymbol{\Sigma} \sim \text{IW}_p(\nu_0, \boldsymbol{\Psi}_0)$, we have $\lambda_1^{p/2}(\boldsymbol{\Psi}_0^{-1}) E\{\text{trace}(\boldsymbol{\Sigma}^{-1})\}^{p/2} \leq E\{\text{trace}(\boldsymbol{\Psi}_0^{-1}\boldsymbol{\Sigma}^{-1})\}^{p/2} < \infty$. Hence, $\int |\boldsymbol{\Sigma}|^{-1/2} \, dP_{0\Sigma}(\boldsymbol{\Sigma}) = \int \prod_{j=1}^p \lambda_j^{1/2}(\boldsymbol{\Sigma}^{-1}) dP_{0\Sigma}(\boldsymbol{\Sigma}) \leq \int \lambda_p^{p/2}(\boldsymbol{\Sigma}^{-1}) dP_{0\Sigma}(\boldsymbol{\Sigma}) = E\lambda_p^{p/2}(\boldsymbol{\Sigma}^{-1}) \leq E\{\text{trace}(\boldsymbol{\Sigma}^{-1})\}^{p/2} < \infty$.

For any two positive semidefinite matrix $\mathbf{A}$ and $\mathbf{B}$, we have $|\mathbf{A}+\mathbf{B}| \geq |\mathbf{A}|$. Therefore, when $\boldsymbol{\Sigma} = \boldsymbol{\Omega} + \boldsymbol{\Lambda}\boldsymbol{\Lambda}^{\mathrm{T}}$, we have $\int |\boldsymbol{\Sigma}|^{-1/2} dP_{0\Sigma}(\boldsymbol{\Sigma}_{1:K}) \leq \int |\boldsymbol{\Omega}|^{-1/2} dP_{0\Sigma}(\boldsymbol{\Sigma}_{1:K}) = \int \prod_{j=1}^{p} \sigma_j^{-1} dP_{0\Sigma}(\boldsymbol{\Sigma}_{1:K}) = \prod_{j=1}^{p} E\sigma_j^{-1} < \infty$, whenever $\sigma_j^2 \sim \text{Inv-Ga}(a,b)$ independently.

The following lemma proves a property of $f_{\boldsymbol{\epsilon}} = \int\int f_{c\boldsymbol{\epsilon}}(\boldsymbol{\epsilon} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}) dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) dP_0(K)$. Here $P_0(K)$ denotes the prior on K, the number of mixture components.

Lemma 8. *Let $f_{0\boldsymbol{\epsilon}} \in \widetilde{\mathcal{F}}_{\boldsymbol{\epsilon}}$ and $f_{\boldsymbol{\epsilon}} \sim \Pi_{\boldsymbol{\epsilon}}$ and $\mathbf{D}(\boldsymbol{\tau}) = diag(\tau_1, \tau_2, \ldots, \tau_p)$. Then*

$$\lim_{\boldsymbol{\tau} \to \mathbf{1}} \int f_{0\boldsymbol{\epsilon}}(\boldsymbol{\epsilon}) \; log \left[\frac{f_{\boldsymbol{\epsilon}}(\boldsymbol{\epsilon})}{|\mathbf{D}(\boldsymbol{\tau})|^{-1} f_{\boldsymbol{\epsilon}}\{\mathbf{D}(\boldsymbol{\tau})\boldsymbol{\epsilon}\}} \right] d\boldsymbol{\epsilon} = 0.$$

Proof. We have $|\mathbf{D}(\boldsymbol{\tau})|^{-1} f_{c\boldsymbol{\epsilon}}\{\mathbf{D}(\boldsymbol{\tau})\boldsymbol{\epsilon}\} \to f_{c\boldsymbol{\epsilon}}(\boldsymbol{\epsilon})$ as $\boldsymbol{\tau} \to \mathbf{1}$. Since $\boldsymbol{\tau} \to \mathbf{1}$, without loss of generality, we may assume $|\mathbf{D}(\boldsymbol{\tau})| > 1/2$. Define $c = \int |\boldsymbol{\Sigma}|^{-1/2} dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Then $c < \infty$. Also $\int |\mathbf{D}(\boldsymbol{\tau})|^{-1} f_{c\boldsymbol{\epsilon}}\{\mathbf{D}(\boldsymbol{\tau})\boldsymbol{\epsilon} \mid \boldsymbol{\theta}\} dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \leq \int 2(2\pi)^{-p/2} |\boldsymbol{\Sigma}|^{-1/2} dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) < 2c < \infty$. Applying DCT, $|\mathbf{D}(\boldsymbol{\tau})|^{-1} f_{\boldsymbol{\epsilon}}\{\mathbf{D}(\boldsymbol{\tau})\boldsymbol{\epsilon}\} \to f_{\boldsymbol{\epsilon}}(\boldsymbol{\epsilon})$ as $\boldsymbol{\tau} \to \mathbf{1}$. Therefore, for any $\boldsymbol{\epsilon} \in \mathbb{R}$,

$$\log \left[\frac{f_{\boldsymbol{\epsilon}}(\boldsymbol{\epsilon})}{|\mathbf{D}(\boldsymbol{\tau})|^{-1} f_{\boldsymbol{\epsilon}}\{\mathbf{D}(\boldsymbol{\tau})\boldsymbol{\epsilon}\}} \right] \to 0 \quad \text{as } \boldsymbol{\tau} \to \mathbf{1}.$$

To find an integrable with respect to $f_{0\boldsymbol{\epsilon}}$ upper bound for $\log\left[|\mathbf{D}(\boldsymbol{\tau})|\, f_{\boldsymbol{\epsilon}}(\boldsymbol{\epsilon}) / f_{\boldsymbol{\epsilon}}\{\mathbf{D}(\boldsymbol{\tau})\boldsymbol{\epsilon}\}\right]$, we use Lemma 7. To do so, we can ignore the prior $P_0(K)$ since the upper bounds obtained in Lemma 7 do not depend on the specific choice of K. We have, using part 3 of Lemma 7,

$$\int |\boldsymbol{\Sigma}|^{-1/2} \exp\left[-\frac{1}{2} \{\mathbf{D}(\boldsymbol{\tau})\boldsymbol{\epsilon} - \boldsymbol{\mu}\}^{\mathrm{T}} \boldsymbol{\Sigma}^{-1} \{\mathbf{D}(\boldsymbol{\tau})\boldsymbol{\epsilon} - \boldsymbol{\mu}\} \right] dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \leq \int |\boldsymbol{\Sigma}|^{-1/2} dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \leq c.$$

Since $\boldsymbol{\tau} \to \mathbf{1}$, without loss of generality we may also assume $\tau_k < 2$ for all k. Therefore,

$$\begin{aligned}
&|\log f_{\boldsymbol{\epsilon}}\{\mathbf{D}(\boldsymbol{\tau})\boldsymbol{\epsilon}\}| \\
&\quad \leq \log(2\pi)^{p/2} + \left| \log \int |\boldsymbol{\Sigma}|^{-1/2} \exp\left[-\frac{1}{2} \{\mathbf{D}(\boldsymbol{\tau})\boldsymbol{\epsilon} - \boldsymbol{\mu}\}^{\mathrm{T}} \boldsymbol{\Sigma}^{-1} \{\mathbf{D}(\boldsymbol{\tau})\boldsymbol{\epsilon} - \boldsymbol{\mu}\} \right] dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \right| \\
&\quad \leq \log(2\pi)^{p/2} + |\log c| \\
&\qquad\qquad - \log \int c^{-1} |\boldsymbol{\Sigma}|^{-1/2} \exp\left[-\frac{1}{2} \{\mathbf{D}(\boldsymbol{\tau})\boldsymbol{\epsilon} - \boldsymbol{\mu}\}^{\mathrm{T}} \boldsymbol{\Sigma}^{-1} \{\mathbf{D}(\boldsymbol{\tau})\boldsymbol{\epsilon} - \boldsymbol{\mu}\} \right] dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \\
&\quad \leq \log\{c(2\pi)^{p/2}\} + |\log c| \\
&\qquad\qquad + \frac{1}{2} \int \log |\boldsymbol{\Sigma}|\, dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) + \frac{1}{2} \int \{\mathbf{D}(\boldsymbol{\tau})\boldsymbol{\epsilon} - \boldsymbol{\mu}\}^{\mathrm{T}} \boldsymbol{\Sigma}^{-1} \{\mathbf{D}(\boldsymbol{\tau})\boldsymbol{\epsilon} - \boldsymbol{\mu}\} dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \\
&\quad \leq \log\{c(2\pi)^{p/2}\} + |\log c| \\
&\qquad\qquad + \frac{1}{2} \int \log |\boldsymbol{\Sigma}|\, dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) + \frac{1}{2} \int \|\mathbf{D}(\boldsymbol{\tau})\boldsymbol{\epsilon} - \boldsymbol{\mu}\|_2^2 \, \lambda_1^{-1}(\boldsymbol{\Sigma}) dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \\
&\quad \leq \log\{c(2\pi)^{p/2}\} + |\log c|
\end{aligned}$$

$$+\frac{1}{2}\int \log|\boldsymbol{\Sigma}|\, dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu},\boldsymbol{\Sigma}) + \int \{\|\mathbf{D}(\boldsymbol{\tau})\boldsymbol{\epsilon}\|_2^2 + \|\boldsymbol{\mu}\|_2^2\}\lambda_1^{-1}(\boldsymbol{\Sigma}) dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu},\boldsymbol{\Sigma})$$

$$\leq \log\{c(2\pi)^{p/2}\} + |\log\, c| + \frac{1}{2}\int \log|\boldsymbol{\Sigma}|\, dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu},\boldsymbol{\Sigma})$$

$$+ \|2\boldsymbol{\epsilon}\|_2^2 \int \lambda_1^{-1}(\boldsymbol{\Sigma}) dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu},\boldsymbol{\Sigma}) + \int \|\boldsymbol{\mu}\|_2^2\, dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu},\boldsymbol{\Sigma}) \int \lambda_1^{-1}(\boldsymbol{\Sigma}) dP_{\boldsymbol{\epsilon},K}(\boldsymbol{\mu},\boldsymbol{\Sigma}),$$

where the third step followed from application of Jensen's inequality on $g(Z) = -\log Z$. The regularity assumptions on $f_{0\boldsymbol{\epsilon}}$ and Lemma 7 imply that the RHS above is $f_{0\boldsymbol{\epsilon}}$ integrable. The conclusion of Lemma 8 follows from an application of DCT again.

To prove Lemma 3, let $f_{\mathbf{U}|\mathbf{S}}$ denote the density of $\mathbf{U} = \mathbf{S}(\mathbf{X})\boldsymbol{\epsilon}$, where $\mathbf{S} = \text{diag}(s_1, \ldots, s_p)$. Then $f_{\mathbf{U}|\mathbf{X}} = f_{\mathbf{U}|\mathbf{S}(\mathbf{X})}$. We have $f_{\mathbf{U}|\mathbf{S}}(\mathbf{U}) = |\mathbf{S}|^{-1} f_{\boldsymbol{\epsilon}}(\mathbf{S}^{-1}\mathbf{U})$. This implies

$$\int f_{0\mathbf{U}|\mathbf{S}_0}(\mathbf{U}) \log\frac{f_{0\mathbf{U}|\mathbf{S}_0}(\mathbf{U})}{f_{\mathbf{U}|\mathbf{S}}(\mathbf{U})} d\mathbf{U} = \int f_{0\mathbf{U}|\mathbf{S}_0}(\mathbf{U}) \log\frac{f_{0\mathbf{U}|\mathbf{S}_0}(\mathbf{U})}{f_{\mathbf{U}|\mathbf{S}_0}(\mathbf{U})} d\mathbf{U} + \int f_{0\mathbf{U}|\mathbf{S}_0}(\mathbf{U}) \log\frac{f_{\mathbf{U}|\mathbf{S}_0}(\mathbf{U})}{f_{\mathbf{U}|\mathbf{S}}(\mathbf{U})} d\mathbf{U}$$

$$= \int f_{0\boldsymbol{\epsilon}}(\boldsymbol{\epsilon}) \log\frac{f_{0\boldsymbol{\epsilon}}(\boldsymbol{\epsilon})}{f_{\boldsymbol{\epsilon}}(\boldsymbol{\epsilon})} d\boldsymbol{\epsilon} + \int f_{0\boldsymbol{\epsilon}}(\boldsymbol{\epsilon}) \log\frac{f_{\boldsymbol{\epsilon}}(\boldsymbol{\epsilon})}{|\mathbf{S}|^{-1}|\mathbf{S}_0| f_{\boldsymbol{\epsilon}}(\mathbf{S}^{-1}\mathbf{S}_0\boldsymbol{\epsilon})} d\boldsymbol{\epsilon}.$$

Let $\delta > 0$ be given. By part 2 of Lemma 1, $\Pi_{\boldsymbol{\epsilon}}\{f_{\boldsymbol{\epsilon}} : d_{KL}(f_{0\boldsymbol{\epsilon}}, f_{\boldsymbol{\epsilon}}) < \delta/2\} > 0$. Let $\mathbf{s} = (s_1, \ldots, s_p)^{\mathrm{T}}$ and $\mathbf{s}_0 = (s_{01}, \ldots, s_{0p})^{\mathrm{T}}$. By Lemma 8, there exists $\eta > 0$ such that $\|\mathbf{s}_0 - \mathbf{s}\|_\infty < \eta$ implies $\int f_{0\boldsymbol{\epsilon}}(\boldsymbol{\epsilon})\, \log[f_{\boldsymbol{\epsilon}}(\boldsymbol{\epsilon})/\{|\mathbf{S}|^{-1}|\mathbf{S}_0| f_{\boldsymbol{\epsilon}}(\mathbf{S}^{-1}\mathbf{S}_0\boldsymbol{\epsilon})\}]\, d\boldsymbol{\epsilon} < \delta/2$ for every $f_{\boldsymbol{\epsilon}} \sim \Pi_{\boldsymbol{\epsilon}}$. Using a straightforward multivariate extension of Corollary 1, we have $\Pi_{\mathbf{V}}(||\mathbf{s}_0 - \mathbf{s}||_\infty < \eta) > 0$. Combining these results, $\Pi_{\mathbf{U}|\mathbf{V}}\{\sup_{\mathbf{X}\in\mathcal{X}} d_{KL}(f_{0\mathbf{U}|\mathbf{X}}, f_{\mathbf{U}|\mathbf{X}}) < \delta\} \geq \Pi_{\boldsymbol{\epsilon}}\{d_{KL}(f_{0\boldsymbol{\epsilon}}, f_{\boldsymbol{\epsilon}}) < \delta/2\}\ \Pi_{\mathbf{V}}(||\mathbf{s}_0 - \mathbf{s}||_\infty < \eta) > 0$. Hence the proof of part 2 of Lemma 3.

Part 1 of Lemma 3 follows trivially from part 2 of Lemma 3 since $||\mathbf{s}_0 - \mathbf{s}||_\infty < \eta$ implies $\|\mathbf{s}_0(\mathbf{X}) - \mathbf{s}(\mathbf{X})\|_\infty < \eta$ for any $\mathbf{X} \in \mathcal{X}$.

To prove part 3 of Lemma 3, note that

$$d_{KL}(f_{0,\mathbf{X},\mathbf{U}}, f_{\mathbf{X},\mathbf{U}}) = \int_{\mathcal{X}\times\mathbb{R}^p} f_{0,\mathbf{U}|\mathbf{X}}(\mathbf{U} \mid \mathbf{X}) f_{0,\mathbf{X}}(\mathbf{X})\, \log\frac{f_{0,\mathbf{U}|\mathbf{X}}(\mathbf{U} \mid \mathbf{X}) f_{0,\mathbf{X}}(\mathbf{X})}{f_{\mathbf{U}|\mathbf{X}}(\mathbf{U} \mid \mathbf{X}) f_{\mathbf{X}}(\mathbf{X})} d\mathbf{X} d\mathbf{U}$$

$$= \int_{\mathcal{X}} f_{0,\mathbf{X}}(\mathbf{X}) \int_{\mathbb{R}^p} f_{0,\mathbf{U}|\mathbf{X}}(\mathbf{U} \mid \mathbf{X})\, \log\frac{f_{0,\mathbf{U}|\mathbf{X}}(\mathbf{U} \mid \mathbf{X})}{f_{\mathbf{U}|\mathbf{X}}(\mathbf{U} \mid \mathbf{X})} d\mathbf{U} d\mathbf{X} + \int_{\mathcal{X}} f_{0,\mathbf{X}}(\mathbf{X})\, \log\frac{f_{0,\mathbf{X}}(\mathbf{X})}{f_{\mathbf{X}}(\mathbf{X})} d\mathbf{X}$$

$$\leq \sup_{\mathbf{X}\in\mathcal{X}} d_{KL}\{f_{0,\mathbf{U}|\mathbf{X}}(\mathbf{U} \mid \mathbf{X}), f_{\mathbf{U}|\mathbf{X}}(\mathbf{U} \mid \mathbf{X})\} + d_{KL}(f_{0\mathbf{X}}, f_{\mathbf{X}}).$$

Part 3 of Lemma 3 now follows from part 2 of Lemma 3 and part 1 of Lemma 1.

S.4.4 Proof of Theorem 1

Let $d_H(f_0, f) = [\int \{f_0^{1/2}(\mathbf{Z}) - f^{1/2}(\mathbf{Z})\}^2 d\mathbf{Z}]^{1/2}$ denote the Hellinger distance between any two densities f_0 and f. From Chapter 1 of Ghosh and Ramamoorthi (2010), we have

$$d_H^2(f_0, f) \leq ||f_0 - f||_1 \leq 2\, d_{KL}^{1/2}(f_0, f). \tag{S.1}$$

Using (S.1), we have,

$$
\begin{aligned}
||f_{0\mathbf{W}} - f_{\mathbf{W}}||_1 &= \int |f_{0\mathbf{W}}(\mathbf{W}) - f_{\mathbf{W}}(\mathbf{W})| d\mathbf{W} \\
&= \int \left| \int f_{0\mathbf{X}}(\mathbf{X}) f_{0\mathbf{W}|\mathbf{X}}(\mathbf{W}) d\mathbf{X} - \int f_{\mathbf{X}}(\mathbf{X}) f_{\mathbf{W}|\mathbf{X}}(\mathbf{W}) d\mathbf{X} \right| d\mathbf{W} \\
&\leq \int \left| \int f_{0\mathbf{X}}(\mathbf{X}) f_{0\mathbf{W}|\mathbf{X}}(\mathbf{W}) d\mathbf{X} - \int f_{\mathbf{X}}(\mathbf{X}) f_{0\mathbf{W}|\mathbf{X}}(\mathbf{W}) d\mathbf{X} \right| d\mathbf{W} \\
&\quad + \int \left| \int f_{\mathbf{X}}(\mathbf{X}) f_{0\mathbf{W}|\mathbf{X}}(\mathbf{W}) d\mathbf{X} - \int f_{\mathbf{X}}(\mathbf{X}) f_{\mathbf{W}|\mathbf{X}}(\mathbf{W}) d\mathbf{X} \right| d\mathbf{W} \\
&\leq \int \int |f_{0\mathbf{X}}(\mathbf{X}) - f_{\mathbf{X}}(\mathbf{X})| f_{0\mathbf{W}|\mathbf{X}}(\mathbf{W}) d\mathbf{X} d\mathbf{W} \\
&\quad + \int \int f_{\mathbf{X}}(\mathbf{X}) |f_{0\mathbf{W}|\mathbf{X}}(\mathbf{W}) - f_{\mathbf{W}|\mathbf{X}}(\mathbf{W})| d\mathbf{X} d\mathbf{W} \\
&= \int |f_{0\mathbf{X}}(\mathbf{X}) - f_{\mathbf{X}}(\mathbf{X})| d\mathbf{X} + \int f_{\mathbf{X}}(\mathbf{X}) \int |f_{0\mathbf{W}|\mathbf{X}}(\mathbf{W}) - f_{\mathbf{W}|\mathbf{X}}(\mathbf{W})| d\mathbf{W} d\mathbf{X} \\
&= \int |f_{0\mathbf{X}}(\mathbf{X}) - f_{\mathbf{X}}(\mathbf{X})| d\mathbf{X} + \int f_{\mathbf{X}}(\mathbf{X}) \int |f_{0\mathbf{U}|\mathbf{X}}(\mathbf{W} - \mathbf{X}) - f_{\mathbf{U}|\mathbf{X}}(\mathbf{W} - \mathbf{X})| d\mathbf{W} d\mathbf{X} \\
&\leq ||f_{0\mathbf{X}} - f_{\mathbf{X}}||_1 + \sup_{\mathbf{X} \in \mathcal{X}} ||f_{0\mathbf{U}|\mathbf{X}} - f_{\mathbf{U}|\mathbf{X}}||_1 \\
&\leq 2\, d_{KL}^{1/2}(f_{0\mathbf{X}}, f_{\mathbf{X}}) + 2 \sup_{\mathbf{X} \in \mathcal{X}} d_{KL}^{1/2}(f_{0\mathbf{U}|\mathbf{X}}, f_{\mathbf{U}|\mathbf{X}}).
\end{aligned}
$$

The proof of Theorem 1 follows by combining part 1 of Lemma 1 and part 2 of Lemma 3.

S.5 Additional Figures

We first present, in Subsection S.5.1, some additional figures summarizing the results of the simulation experiments for diagonal covariance matrices discussed in Section 7 of the main paper. Then in Subsection S.5.1, we present figures that summarize the results of simulation experiments for covariance matrices with AR structure. Finally in Subsection S.5.3, we present some additional figures summarizing the results of the EATS data set analyzed in Section 8 of the main paper.

S.5.1 Additional Figures Summarizing the Results of the Simulation Experiments for Diagonal Covariance Structure

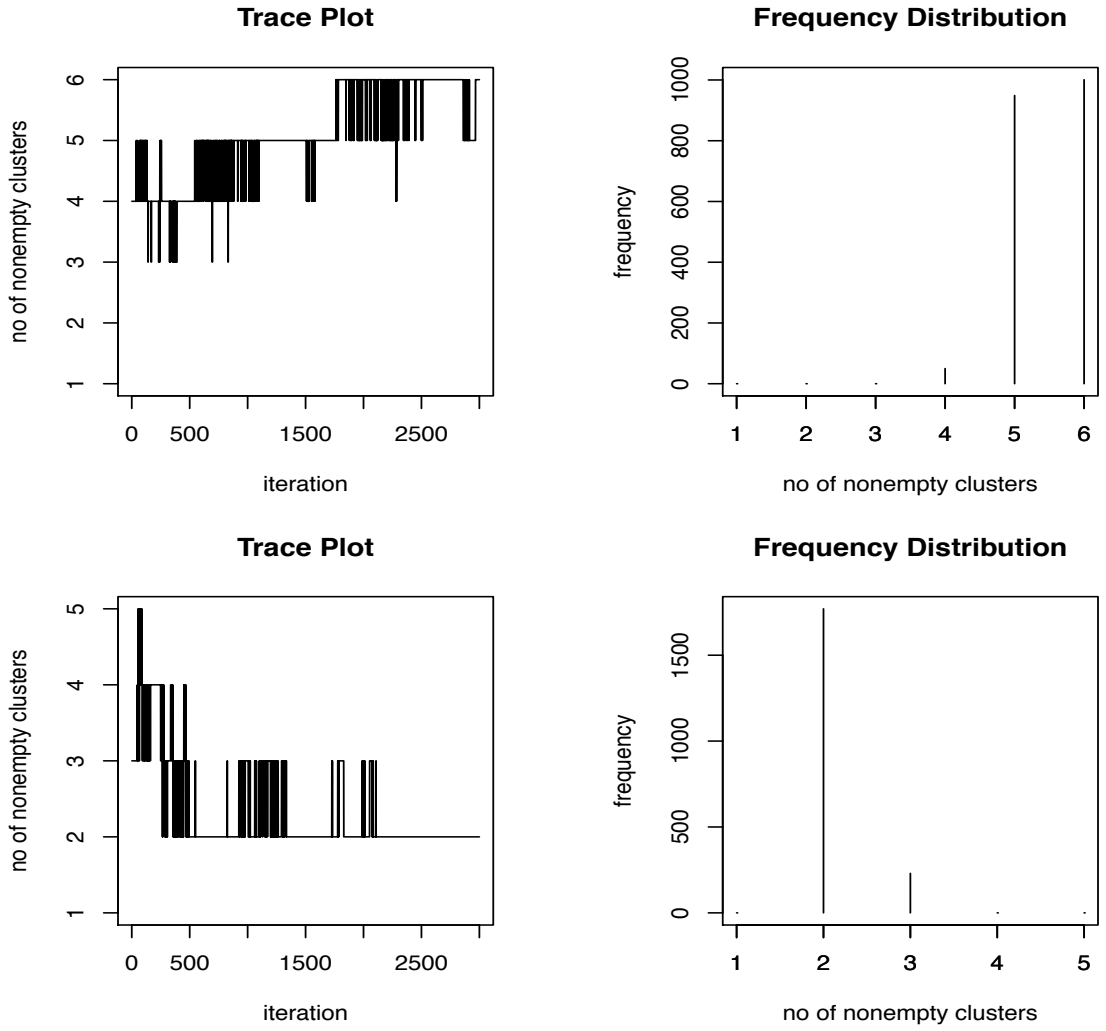

Figure S.2: Trace plots and frequency distributions of the number of nonempty clusters produced by the MIW (mixtures with inverse Wishart priors) method for the conditionally heteroscedastic error distribution $f_{\boldsymbol{\epsilon}}^{(2)}$ with sample size $n = 1000$, $m_i = 3$ replicates for each subject and identity matrix (I) for the component specific covariance matrices. See Section 7 for additional details. The results correspond to the simulation instance that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets, when the number of mixture components for $f_{\mathbf{X}}$ and $f_{\boldsymbol{\epsilon}}$ were kept fixed at $K_{\mathbf{X}} = 6$ and $K_{\boldsymbol{\epsilon}} = 5$. The upper panels are for the $f_{\mathbf{X}}$ and the lower panels are for the density of the scaled errors $f_{\boldsymbol{\epsilon}}$. The true number of mixture components were $K_{\mathbf{X}} = 3$ and $K_{\boldsymbol{\epsilon}} = 3$. As can be seen from Figure 4, a mixture model with 2 nonempty clusters can approximate the true density of the scaled errors well.

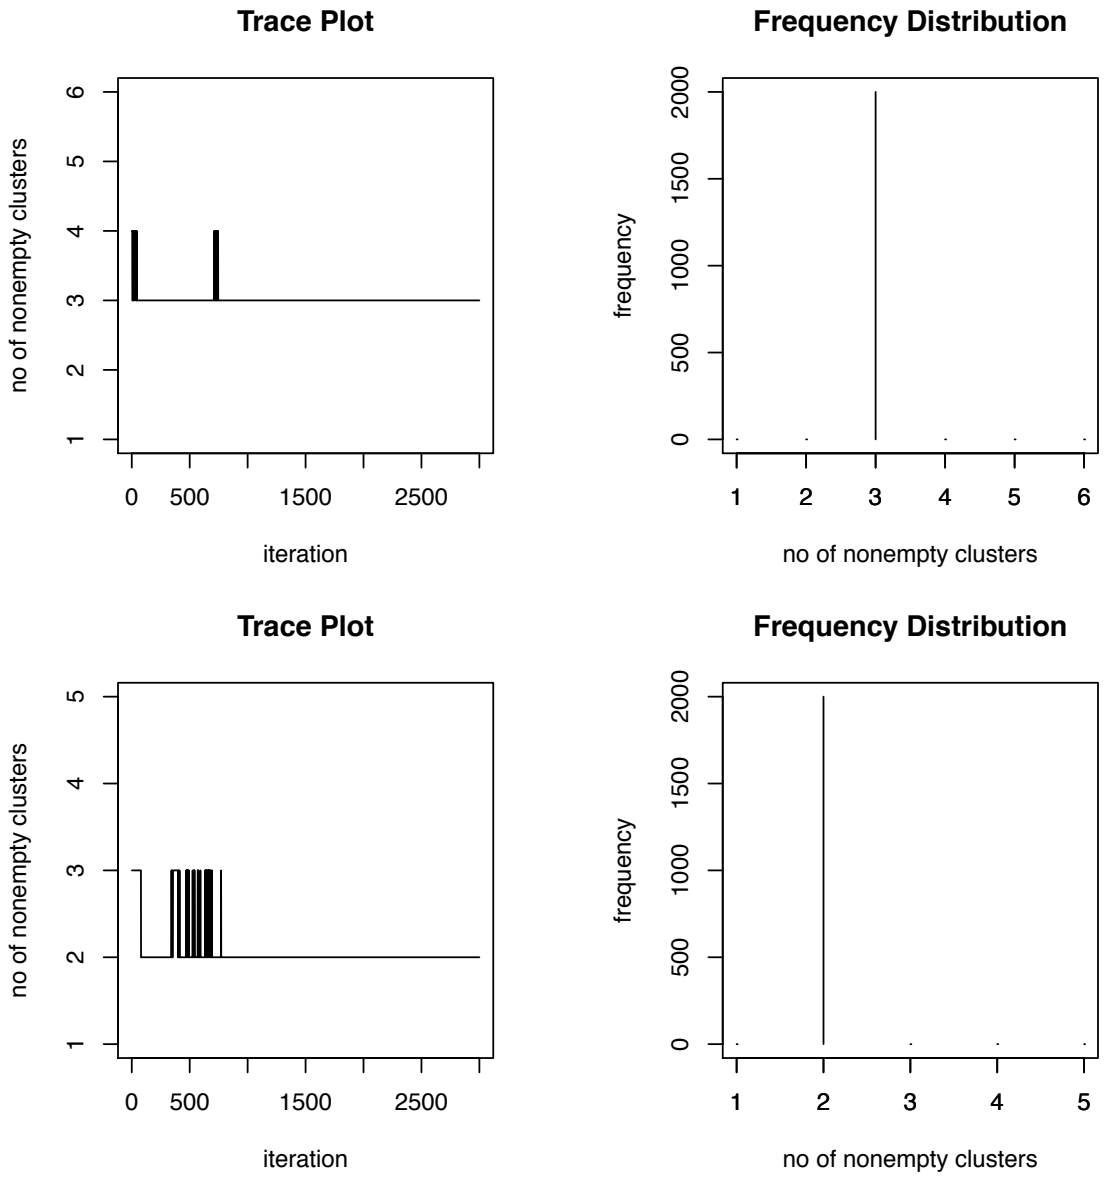

Figure S.3: Trace plots and frequency distributions of the number of nonempty clusters produced by the MLFA (mixtures of latent factor analyzers) method for the conditionally heteroscedastic error distribution $f_{\boldsymbol{\epsilon}}^{(2)}$ with sample size $n = 1000$, $m_i = 3$ replicates for each subject and identity matrix (I) for the component specific covariance matrices. See Section 7 for additional details. The results correspond to the simulation instance that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets, when the number of mixture components for $f_{\mathbf{X}}$ and $f_{\boldsymbol{\epsilon}}$ were kept fixed at $K_{\mathbf{X}} = 6$ and $K_{\boldsymbol{\epsilon}} = 5$. The upper panels are for the $f_{\mathbf{X}}$ and the lower panels are for the density of the scaled errors $f_{\boldsymbol{\epsilon}}$. The true number of mixture components were $K_{\mathbf{X}} = 3$ and $K_{\boldsymbol{\epsilon}} = 3$. As can be seen from Figure 5, a mixture model with 2 nonempty clusters can approximate the true density of the scaled errors well.

S.5.2 Additional Figures Summarizing the Results of the Simulation Experiments for AR Covariance Structure

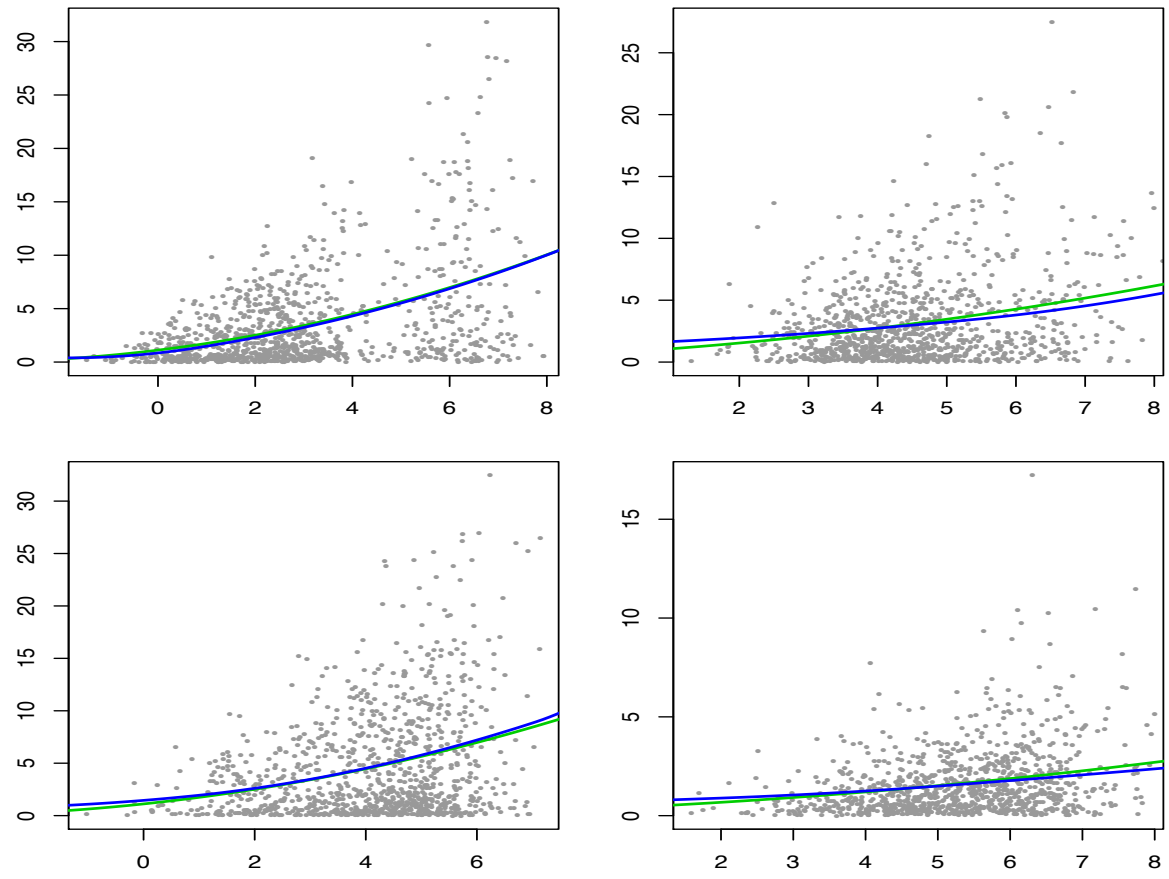

Figure S.4: Results for the variance functions $s^2(X)$ produced by the univariate density deconvolution method for each component of $\mathbf{X}$ for the conditionally heteroscedastic error distribution $f_{\boldsymbol{\epsilon}}^{(2)}$ with sample size $n = 1000$, $m_i = 3$ replicates for each subject and component specific covariance matrices with autoregressive structure (AR). The results correspond to the data set that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets for the MIW (mixtures with inverse Wishart priors) method. For each component of $\mathbf{X}$, the true variance function is $s^2(X) = (1 + X/4)^2$. See Section 2.2.2 and Section 6 for additional details. In each panel, the true (lighter shaded lines) and the estimated (darker shaded lines) variance functions are superimposed over a plot of subject specific sample means vs subject specific sample variances. The figure is in color in the electronic version of this article.

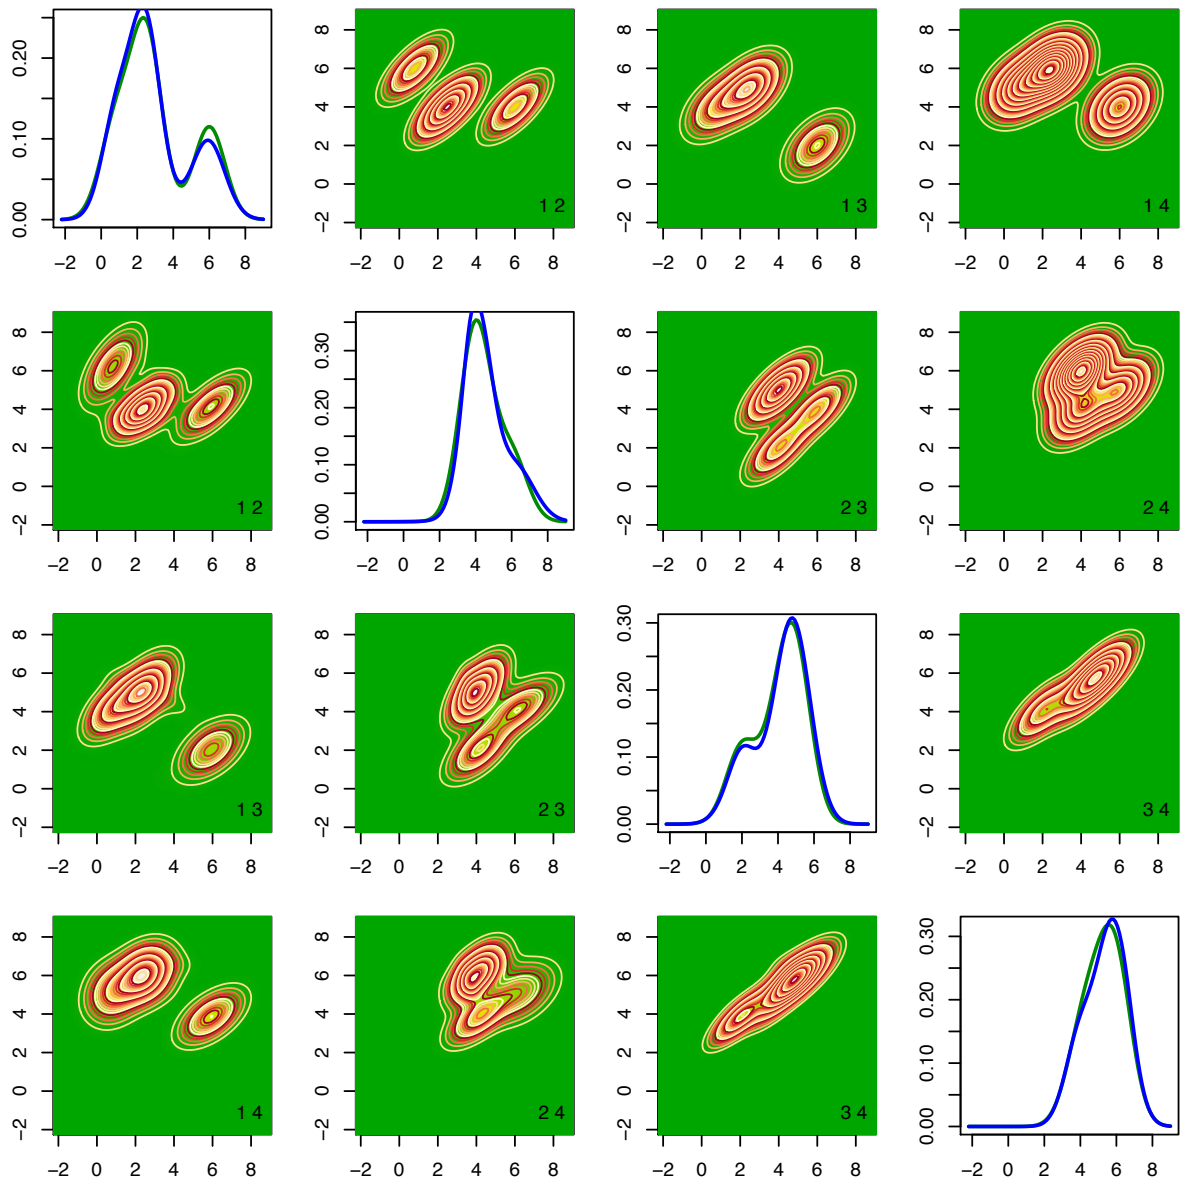

Figure S.5: Results for the $f_{\mathbf{X}}$ produced by the MIW (mixtures with inverse Wishart priors) method for the conditionally heteroscedastic error distribution $f_{\boldsymbol{\epsilon}}^{(2)}$ with sample size $n = 1000$, $m_i = 3$ replicates for each subject and component specific covariance matrices with autoregressive structure (AR). The results correspond to the data set that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets. See Section 7 for additional details. The upper triangular panels show the contour plots of the true two dimensional marginal densities. The lower triangular diagonally opposite panels show the corresponding estimates. The diagonal panels show the true (lighter shaded lines) and the estimated (darker shaded lines) one dimensional marginals. The figure is in color in the electronic version of this article.

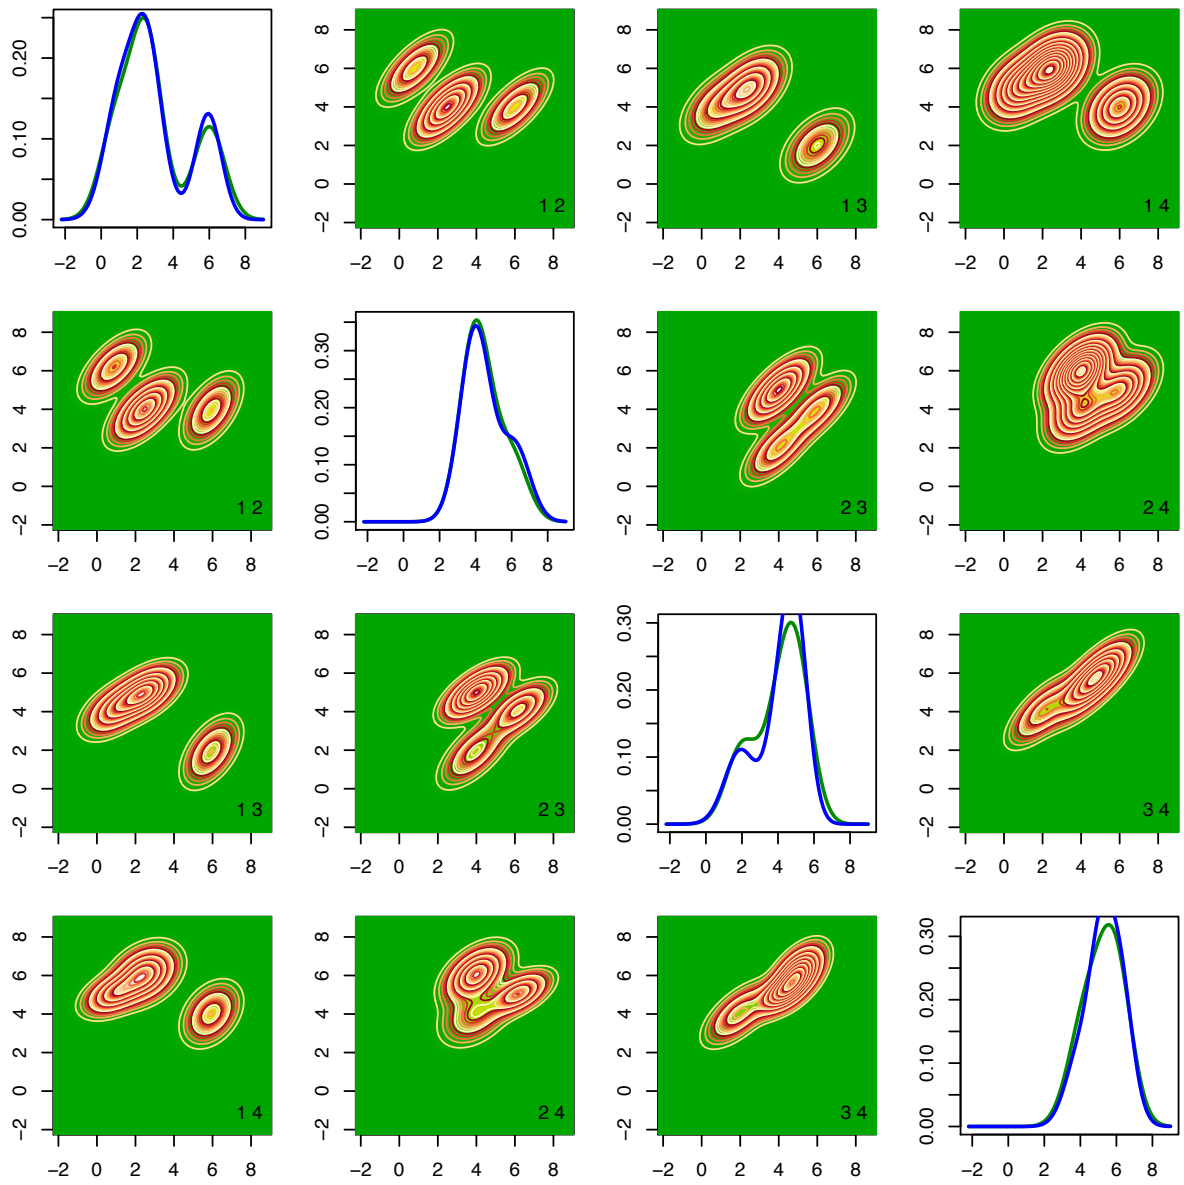

Figure S.6: Results for the $f_{\mathbf{X}}$ produced by the MLFA (mixtures of latent factor analyzers) method for the conditionally heteroscedastic error distribution $f_{\boldsymbol{\epsilon}}^{(2)}$ with sample size $n = 1000$, $m_i = 3$ replicates for each subject and component specific covariance matrices with autoregressive structure (AR). The results correspond to the data set that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets. See Section 7 for additional details. The upper triangular panels show the contour plots of the true two dimensional marginal densities. The lower triangular diagonally opposite panels show the corresponding estimates. The diagonal panels show the true (lighter shaded lines) and the estimated (darker shaded lines) one dimensional marginals. The figure is in color in the electronic version of this article.

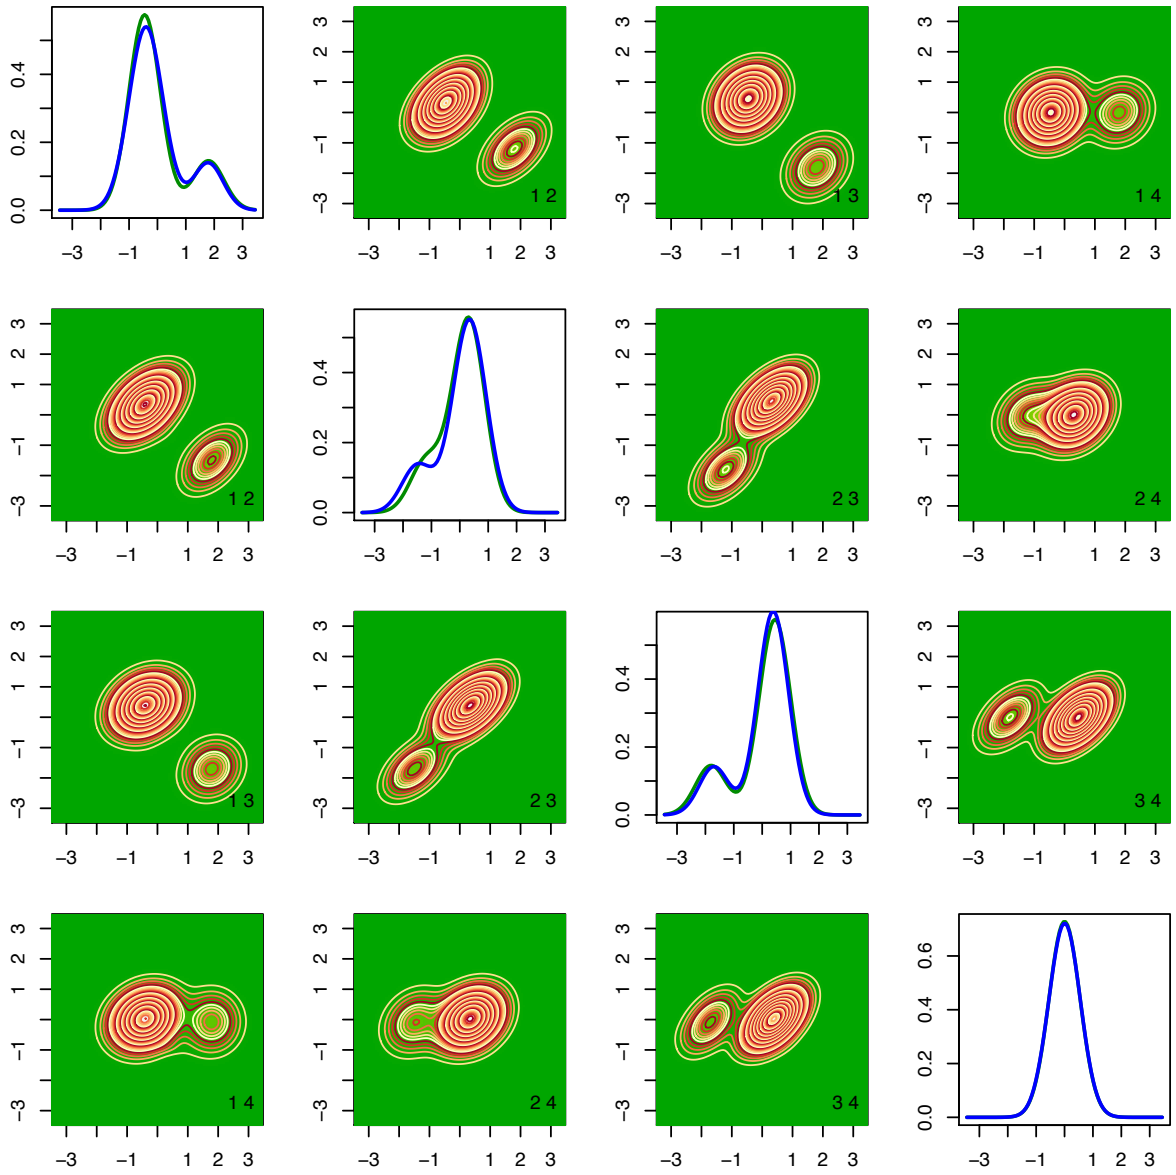

Figure S.7: Results for the density of the scaled errors $f_{\boldsymbol{\epsilon}}$ produced by the MIW (mixtures with inverse Wishart priors) method for the conditionally heteroscedastic error distribution $f_{\boldsymbol{\epsilon}}^{(2)}$ with sample size $n = 1000$, $m_i = 3$ replicates for each subject and component specific covariance matrices with autoregressive structure (AR). The results correspond to the data set that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets. See Section 7 for additional details. The upper triangular panels show the contour plots of the true two dimensional marginal densities. The lower triangular diagonally opposite panels show the corresponding estimates. The diagonal panels show the true (lighter shaded lines) and the estimated (darker shaded lines) one dimensional marginals. The figure is in color in the electronic version of this article.

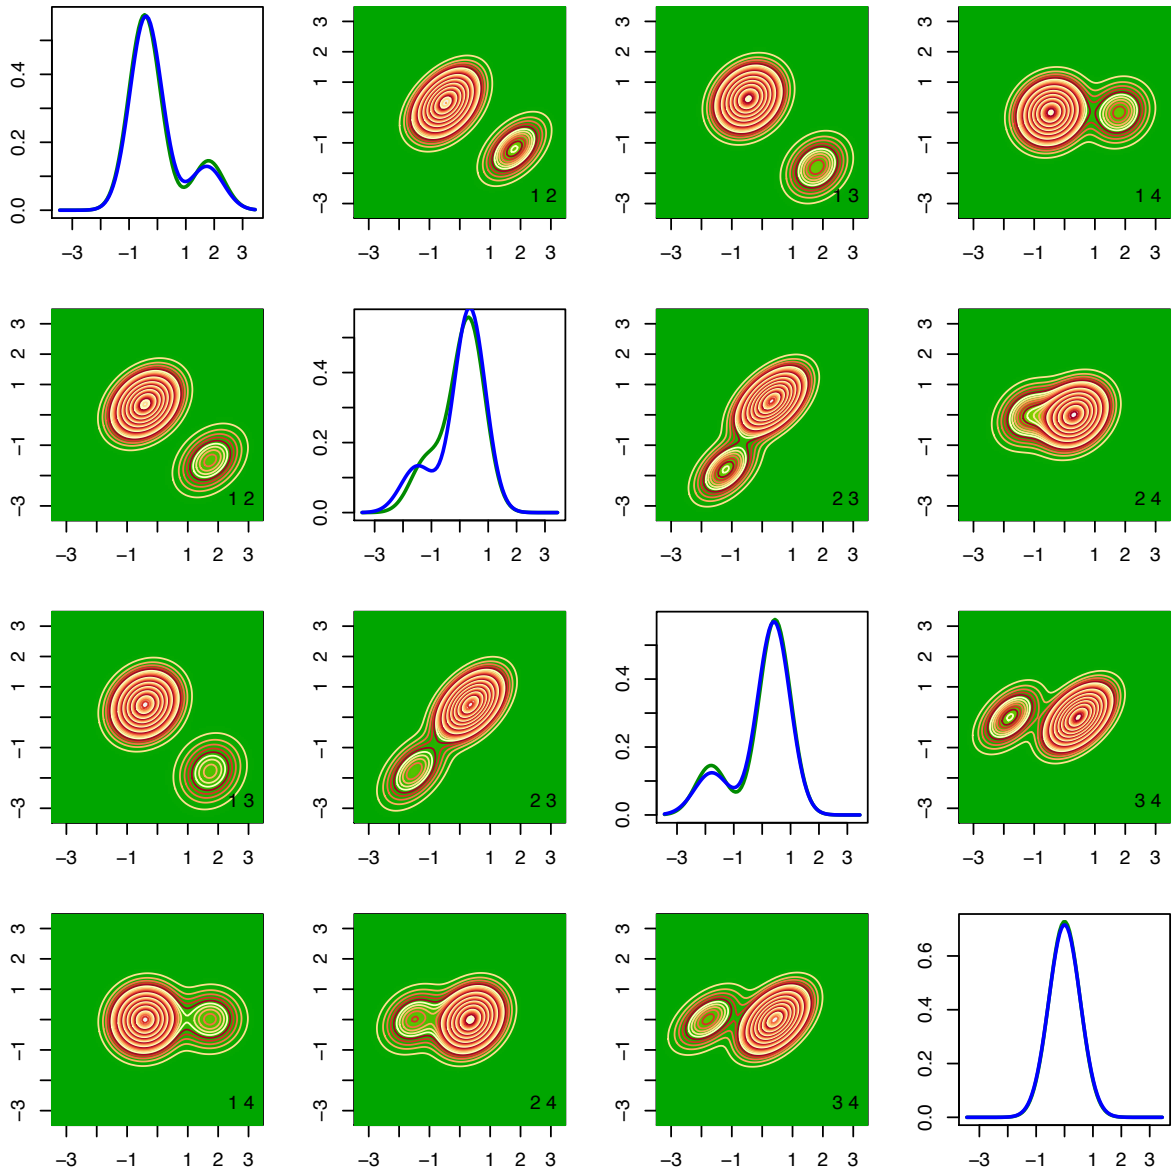

Figure S.8: Results for the density of the scaled errors $f_{\boldsymbol{\epsilon}}$ produced by the MLFA (mixtures of latent factor analyzers) method for the conditionally heteroscedastic error distribution $f_{\boldsymbol{\epsilon}}^{(2)}$ with sample size $n = 1000$, $m_i = 3$ replicates for each subject and component specific covariance matrices with autoregressive structure (AR). The results correspond to the data set that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets. See Section 7 for additional details. The upper triangular panels show the contour plots of the true two dimensional marginal densities. The lower triangular diagonally opposite panels show the corresponding estimates. The diagonal panels show the true (lighter shaded lines) and the estimated (darker shaded lines) one dimensional marginals. The figure is in color in the electronic version of this article.

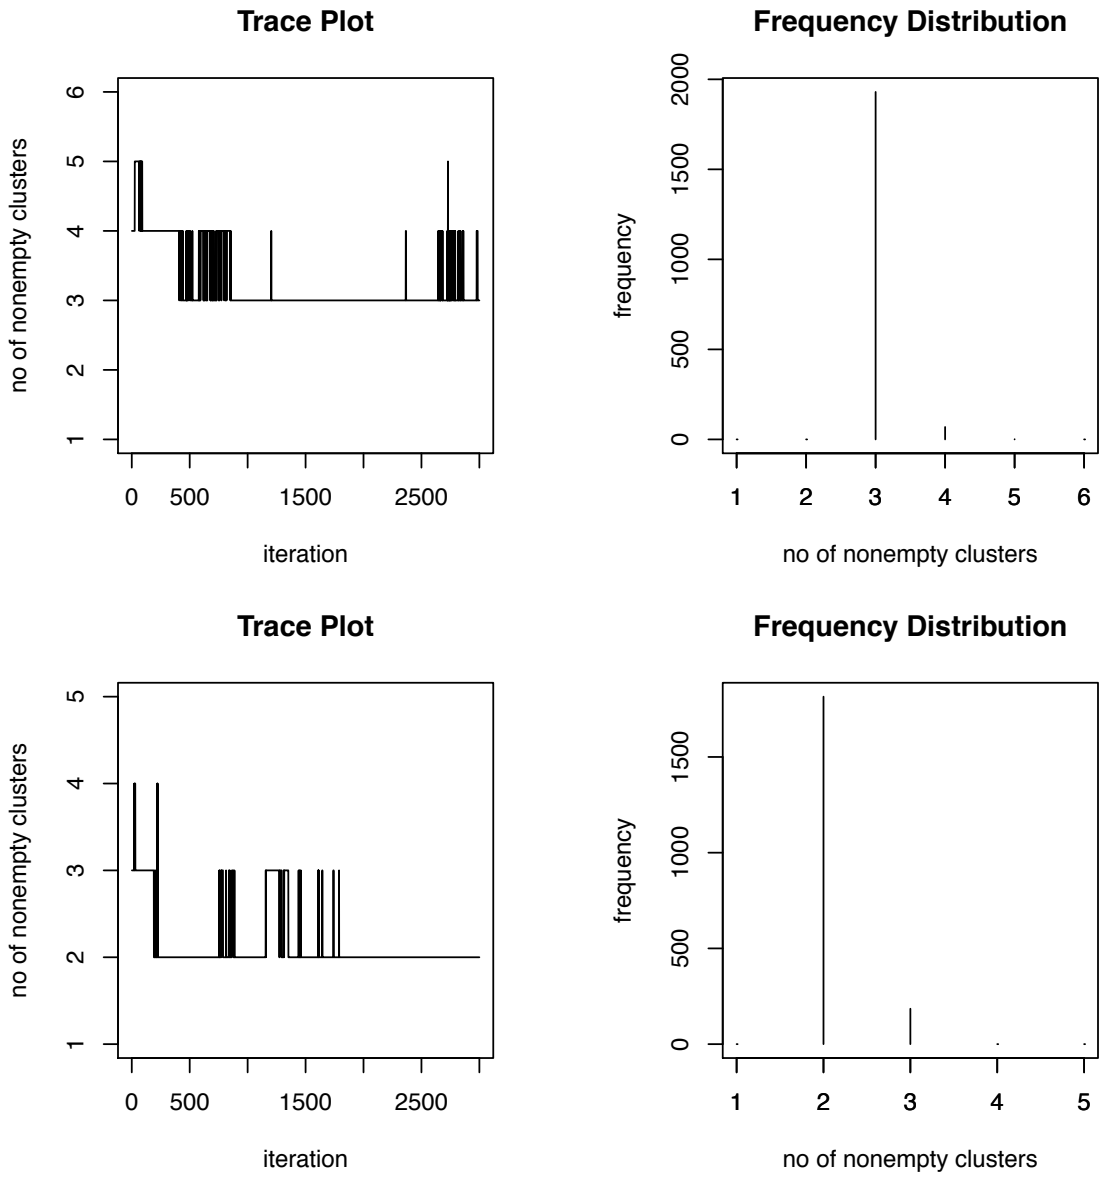

Figure S.9: Trace plots and frequency distributions of the number of nonempty clusters produced by the MIW (mixtures with inverse Wishart priors) method for the conditionally heteroscedastic error distribution $f_{\boldsymbol{\epsilon}}^{(2)}$ with sample size $n = 1000$, $m_i = 3$ replicates for each subject and component specific covariance matrices with autoregressive structure (AR). See Section 7 for additional details. The results correspond to the simulation instance that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets, when the number of mixture components for both $f_{\mathbf{X}}$ and $f_{\boldsymbol{\epsilon}}$ were kept fixed at $K_{\mathbf{X}} = 6$ and $K_{\boldsymbol{\epsilon}} = 5$. The upper panels are for the $f_{\mathbf{X}}$ and the lower panels are for the density of the scaled errors $f_{\boldsymbol{\epsilon}}$. The true number of mixture components were $K_{\mathbf{X}} = 3$ and $K_{\boldsymbol{\epsilon}} = 3$. As can be seen from Figure S.16, a mixture model with 2 nonempty clusters can approximate the true density of the scaled errors well.

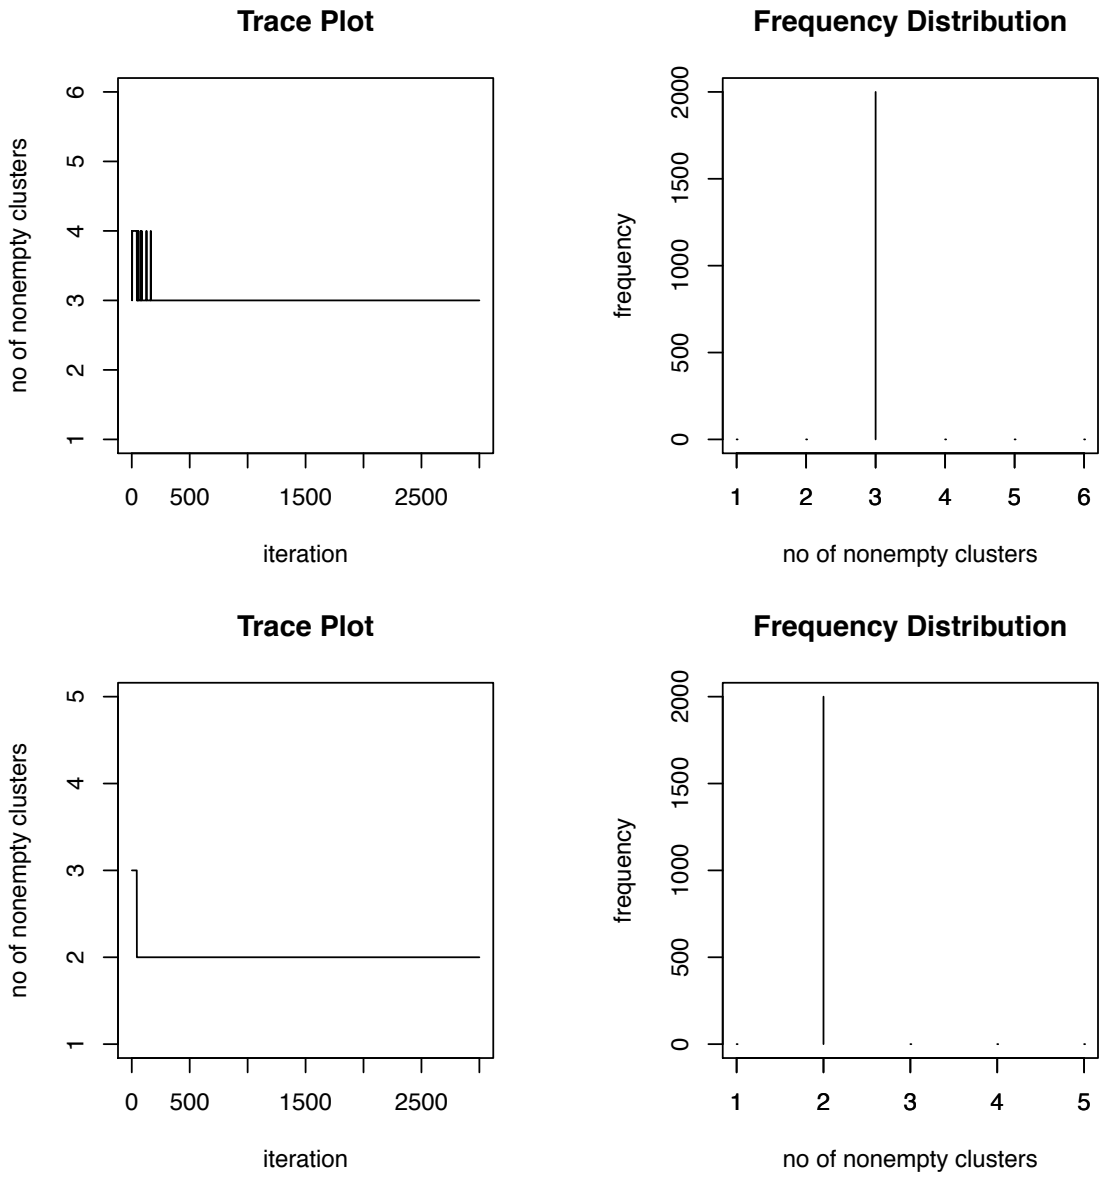

Figure S.10: Trace plots and frequency distributions of the number of nonempty clusters produced by the MLFA (mixtures of latent factor analyzers) method for the conditionally heteroscedastic error distribution $f_{\boldsymbol{\epsilon}}^{(2)}$ with sample size $n = 1000$, $m_i = 3$ replicates for each subject and component specific covariance matrices with autoregressive structure (AR). See Section 7 for additional details. The results correspond to the simulation instance that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets, when the number of mixture components for $f_{\mathbf{X}}$ and $f_{\boldsymbol{\epsilon}}$ were kept fixed at $K_{\mathbf{X}} = 6$ and $K_{\boldsymbol{\epsilon}} = 5$. The upper panels are for the $f_{\mathbf{X}}$ and the lower panels are for the density of the scaled errors $f_{\boldsymbol{\epsilon}}$. The true number of mixture components were $K_{\mathbf{X}} = 3$ and $K_{\boldsymbol{\epsilon}} = 3$. As can be seen from Figure S.8, a mixture model with 2 nonempty clusters can approximate the true density of the scaled errors well.

S.5.3 Additional Figures Summarizing the Results for the EATS Data Set Analyzed in Section 8 of the Main Paper

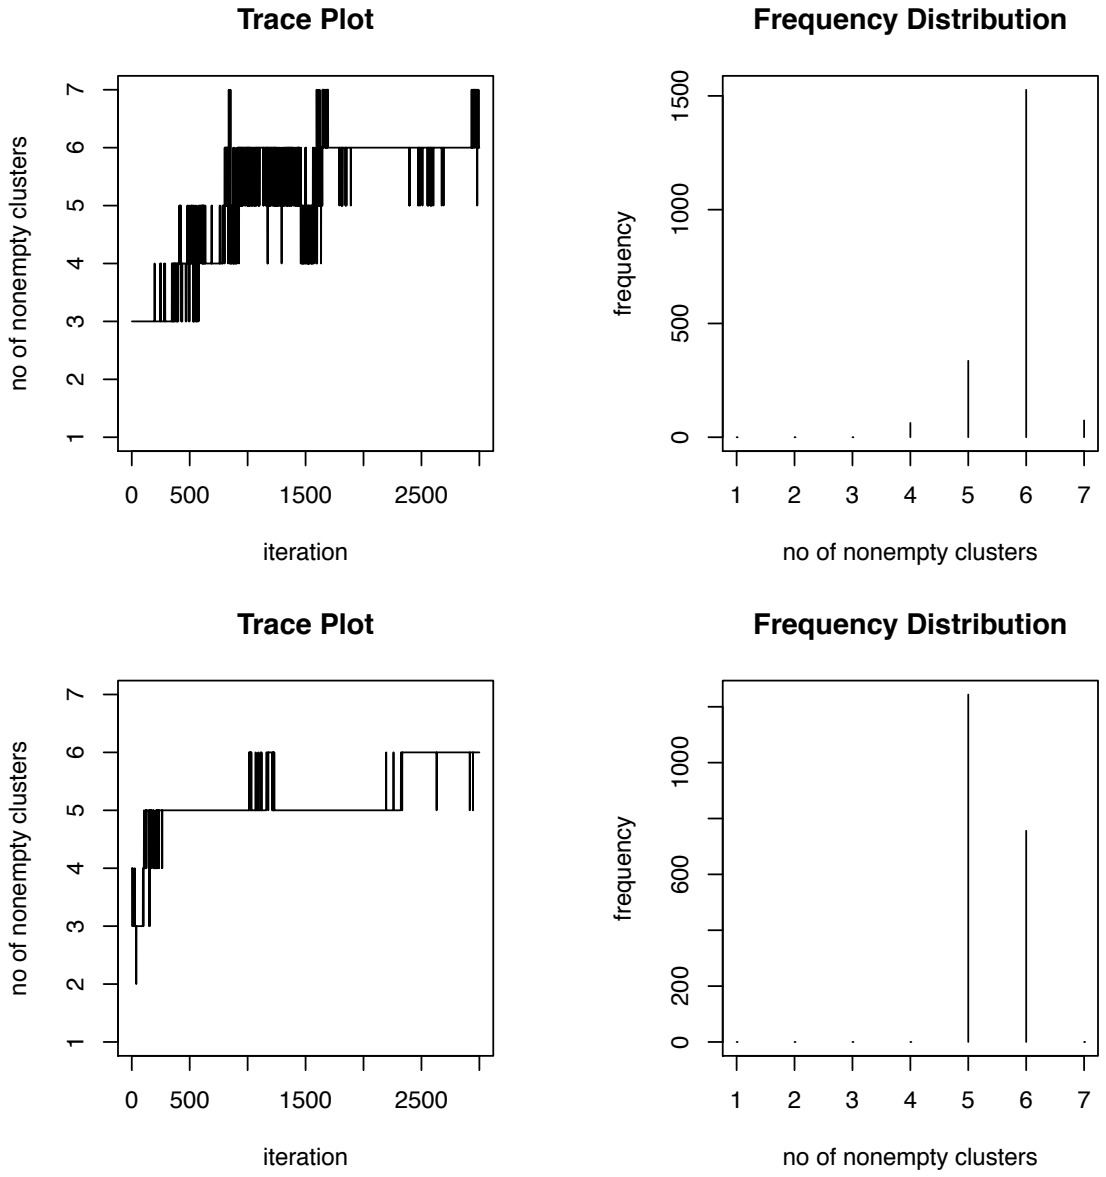

Figure S.11: Trace plots and frequency distributions of the number of nonempty clusters produced by the MIW (mixtures with inverse Wishart priors) method for the EATS data example. See Section 8 for additional details. The number of mixture components for both $f_{\mathbf{X}}$ and $f_{\boldsymbol{\epsilon}}$ were kept fixed at $K_{\mathbf{X}} = K_{\boldsymbol{\epsilon}} = 7$. The upper panels are for the $f_{\mathbf{X}}$ and the lower panels are for the density of the scaled errors $f_{\boldsymbol{\epsilon}}$.

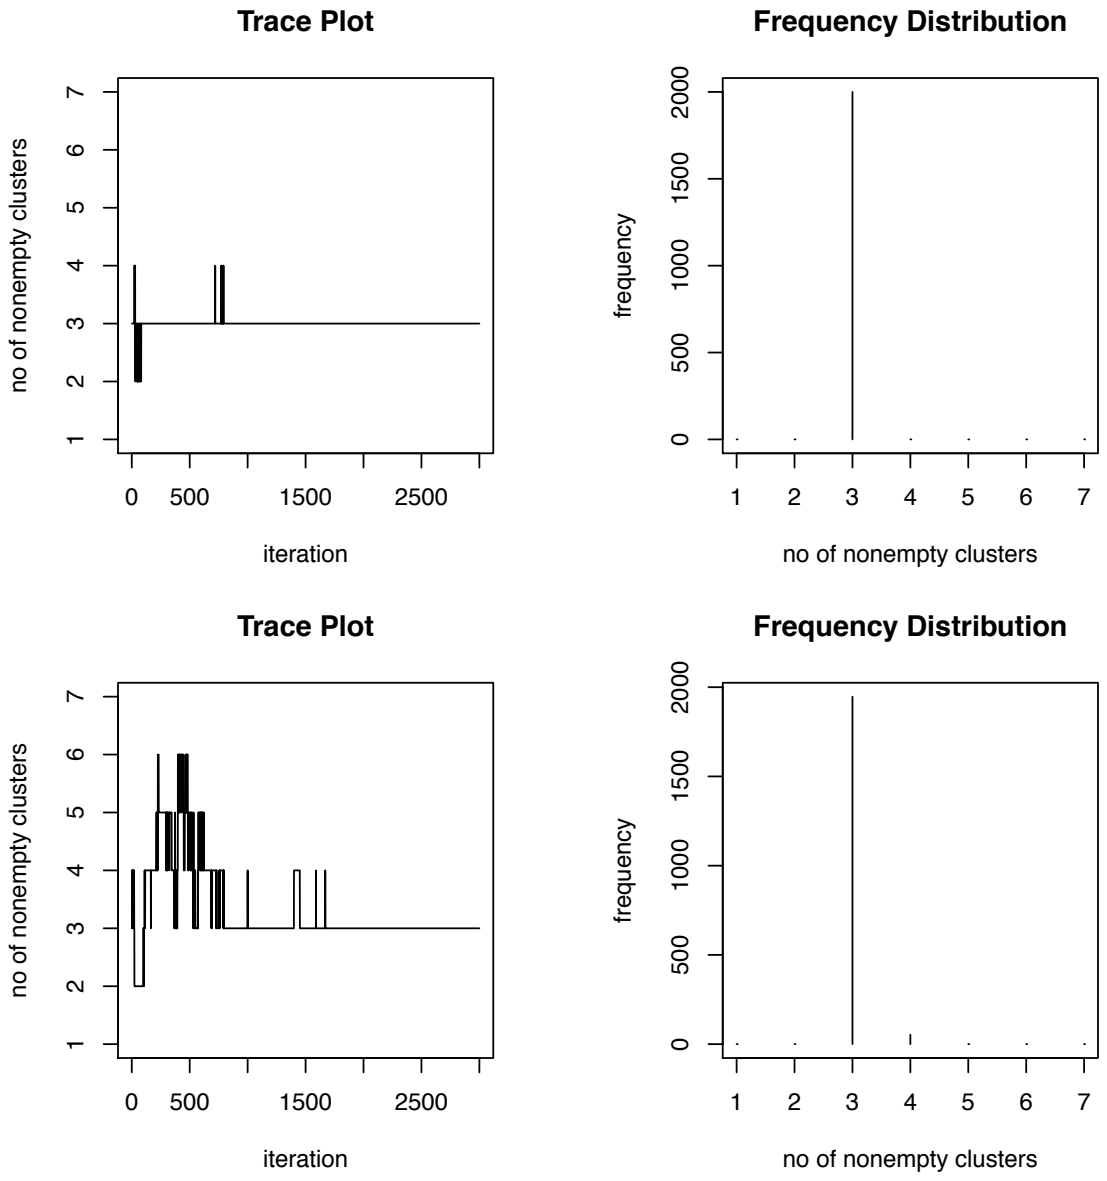

Figure S.12: Trace plots and frequency distributions of the number of nonempty clusters produced by the MLFA (mixtures of latent factor analyzers) method for the EATS data example. See Section 8 for additional details. The number of mixture components for both $f_{\mathbf{X}}$ and $f_{\boldsymbol{\epsilon}}$ were kept fixed at $K_{\mathbf{X}} = K_{\boldsymbol{\epsilon}} = 7$. The upper panels are for the $f_{\mathbf{X}}$ and the lower panels are for the density of the scaled errors $f_{\boldsymbol{\epsilon}}$.

S.6 Additional Simulation Experiments

This section presents the results of additional simulation experiments for multivariate t and multivariate Laplace distributed measurement errors. For easy reference, we provide brief descriptions of the multivariate t and multivariate Laplace distributions below.

S.6.1 Multivariate t Errors

A random variable Z following a Student's t-distribution with degrees of freedom ν, mean μ and variance $\nu b/(\nu-2)$ can be represented as $Z = \mu + \nu^{1/2} b^{1/2} X/Y^{1/2}$, where Y and X are independent, Y follows a chi-square distribution with ν degrees of freedom, denoted by $Y \sim \chi^2_\nu$, and X follows a standard normal distribution. A natural extension to multivariate set up is given by $\mathbf{Z} = \boldsymbol{\mu} + \nu^{1/2}\boldsymbol{\Sigma}^{1/2}\mathbf{X}/Y^{1/2}$, where $Y \sim \chi^2_\nu$ and $\mathbf{X} \sim \text{MVN}_p(\mathbf{0}, \mathbf{I})$ independently. The random vector $\mathbf{Z}$ is then said to follow a multivariate t-distribution (Kotz and Nadarajah, 2004) with degrees of freedom ν, mean $\boldsymbol{\mu}$ and covariance $\nu\boldsymbol{\Sigma}/(\nu-2)$, denoted by $\text{MVT}_p(\nu, \boldsymbol{\mu}, \boldsymbol{\Sigma})$. The above characterization can be used to sample from a $\text{MVT}_p(\nu, \boldsymbol{\mu}, \boldsymbol{\Sigma})$ density. The density of $\mathbf{Z}$ is given by

$$f_{\mathbf{Z}}(\mathbf{z}) = \frac{\Gamma\{(\nu+p)/2\}}{\Gamma(\nu/2)(\nu\pi)^{p/2}\,|\boldsymbol{\Sigma}|^{1/2}} \cdot \{1 + (\mathbf{z}-\boldsymbol{\mu})^{\mathrm{T}}\boldsymbol{\Sigma}^{-1}(\mathbf{z}-\boldsymbol{\mu})/\nu\}^{-(\nu+p)/2}.$$

The characteristic function is given by

$$\phi(\mathbf{t}) = \exp(i\mathbf{t}^{\mathrm{T}}\boldsymbol{\mu}) \cdot \frac{||\nu^{1/2}\boldsymbol{\Sigma}^{1/2}\mathbf{t}||^{\nu/2}}{2^{\nu/2-1}\Gamma(\nu/2)} \cdot H_{\nu/2}(||\nu^{1/2}\boldsymbol{\Sigma}^{1/2}\mathbf{t}||), \quad \mathbf{t} \in \mathbb{R}^p,$$

where H_α denotes a McDonald's function of order $\alpha(> 1/2)$ and admits the integral representation

$$H_\alpha(t) = (2/t)^\alpha \cdot \frac{\Gamma(\alpha+1/2)}{\pi^{1/2}} \int_0^\infty (1+u^2)cos(tu)du, \quad t > 0.$$

When $\boldsymbol{\Sigma} = \mathbf{I}$, the identity matrix, the components Z_i and Z_j are uncorrelated, but not statistically independent. With $\boldsymbol{\mu} = (\mu_1 \ldots, \mu_p)^{\mathrm{T}}$ and $\boldsymbol{\Sigma} = ((\sigma_{ij}))$, the i^{th} random variable Z_i marginally follows a univariate Student's t-distribution with degrees of freedom ν, mean μ_i and variance $\nu\sigma_{ii}/(\nu-2)$.

S.6.2 Multivariate Laplace Errors

A random variable Z following a Laplace distribution with mean μ and variance b has the density

$$f_Z(z) = (2b)^{-1/2} \exp(-2^{1/2}b^{-1/2}\,|z-\mu|).$$

Z can be represented as $Z = \mu + Y^{1/2}b^{1/2}X$, where Y and X are independent and follow standard exponential and standard normal distributions, respectively. A natural extension

to multivariate set up is given by $\mathbf{Z} = \boldsymbol{\mu} + Y^{1/2}\boldsymbol{\Sigma}^{1/2}\mathbf{X}$, where Y follows a standard exponential density and $\mathbf{X} \sim \text{MVN}_p(\mathbf{0}, \mathbf{I})$ independently of Y. The random vector $\mathbf{Z}$ is then said to follow a multivariate Laplace distribution (Eltoft, et al., 2006) with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$, denoted by $\text{MVL}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. The above characterization can be used to sample from a $\text{MVL}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ density. The density of $\mathbf{Z}$ is then given by

$$f_{\mathbf{Z}}(\mathbf{z}) = \frac{2}{(2\pi)^{p/2}\,|\boldsymbol{\Sigma}|^{1/2}} \cdot \frac{K_{p/2-1}\{2^{1/2}h^{1/2}(\mathbf{z})\}}{\{h(\mathbf{z})/2\}^{p/4-1/2}},$$

where $h(\mathbf{z}) = (\mathbf{z}-\boldsymbol{\mu})^{\text{T}}\boldsymbol{\Sigma}^{-1}(\mathbf{z}-\boldsymbol{\mu})$ and K_m denotes modified Bessel functions of the second kind of order m. Using asymptotic formula for the Bessel functions, namely $K_m(z) = \{\pi/(2z)\}^{1/2}\exp(-z)$ as $|z| \to \infty$, we have

$$f_{\mathbf{Z}}(\mathbf{z}) \approx \frac{\pi^{1/2}}{(2\pi)^{p/2}\,|\boldsymbol{\Sigma}|^{1/2}} \cdot \frac{2^{(p-1)/4}}{h^{(p-1)/4}(\mathbf{z})} \cdot \exp\{-2^{1/2}h^{1/2}(\mathbf{z})\}.$$

The characteristic function is given by $\phi(\mathbf{t}) = \exp(i\mathbf{t}^{\text{T}}\boldsymbol{\mu})(1 + \mathbf{t}^{\text{T}}\boldsymbol{\Sigma}\mathbf{t}/2)^{-1}$ for $\mathbf{t} \in \mathbb{R}^p$. For $p > 1$, the density has a singularity at $\mathbf{0}$. When $\boldsymbol{\Sigma} = \mathbf{I}$, the identity matrix, the components Z_i and Z_j are uncorrelated, but not statistically independent. With $\boldsymbol{\mu} = (\mu_1 \ldots, \mu_p)^{\text{T}}$ and $\boldsymbol{\Sigma} = ((\sigma_{ij}))$, the i^{th} random variable Z_i marginally follows a univariate Laplace distribution with mean μ_i and variance σ_{ii}.

S.6.3 Summary of Results

The results of the simulation experiments when the measurement errors are distributed according to $f_{\boldsymbol{\epsilon}}^{(4)} = \text{MVT}_4(6, \mathbf{0}, \boldsymbol{\Sigma})$ and $f_{\boldsymbol{\epsilon}}^{(4)} = \text{MVL}_4(\mathbf{0}, \boldsymbol{\Sigma})$ probability laws independently of $\mathbf{X}$ are presented in Table S.1. The results for conditionally heteroscedastic measurement errors are presented in Table S.2. As in the main paper, in each case four different choices for the covariance matrix $\boldsymbol{\Sigma}$ were considered. The general patterns of the estimated MISEs are similar to that observed in Table 2 and Table 2 of the main paper where the true measurement error distributions were finite mixtures of multivariate normal kernels. While in theory the MLFA model described in the main paper can approximate distributions like the multivariate Laplace that puts significant mass around the origin, in practice, since it assumes $\boldsymbol{\Omega}_k = \boldsymbol{\Omega} = \text{diag}\{\sigma_1^2, \ldots, \sigma_p^2\}$ for all k, it often smooths out the spikes at the origin. A mild variation, referred to as the MLFAD model, that instead assumes $\boldsymbol{\Omega}_k = \sigma_k^2\text{I}_p$ and results in slight improvement in the MISE performance is also included in Table S.1 and Table S.2. For the simulation experiments and the real data analysis presented in the main text, the two versions of the MLFA model perform very similarly and the latter version was not included. Results for conditionally heteroscedastic multivariate Laplace errors with diagonal covariance structure are summarized in Figures S.13-S.19 with observations similar to those discussed in Section 7 of the main paper.

True Error Distribution	Covariance Structure	Sample Size	MISE $\times 10^4$			
			MLFAD	MLFA	MIW	Naive
(c) Multivariate t	I	500	**1.06**	1.38	3.98	12.32
		1000	**0.53**	0.65	1.54	9.91
	LF	500	**6.62**	8.26	7.57	47.22
		1000	4.73	5.78	**3.65**	45.70
	AR	500	12.69	13.56	**6.14**	40.76
		1000	11.36	9.16	**3.45**	39.59
	EXP	500	7.84	8.42	**5.00**	26.85
		1000	6.26	6.64	**2.38**	26.04
(d) Multivariate Laplace	I	500	**1.08**	1.32	3.08	8.22
		1000	**0.50**	0.63	1.20	6.25
	LF	500	**4.41**	5.57	5.66	32.31
		1000	**2.38**	3.53	2.84	31.10
	AR	500	8.38	8.72	**5.14**	27.30
		1000	6.08	6.19	**2.56**	26.19
	EXP	500	5.24	5.67	**4.14**	17.57
		1000	3.58	4.17	**1.98**	16.86

Table S.1: Mean integrated squared error (MISE) performance of MLFA (mixtures of latent factor analyzers) and MIW (mixtures with inverse Wishart priors) density deconvolution models for **homoscedastic** errors compared with a naive method that ignores measurement errors for different measurement error distributions. See Section 2 and Section S.6 for additional details. The minimum value in each row is highlighted.

True Error Distribution	Covariance Structure	Sample Size	MISE $\times 10^4$			
			MLFAD	MLFA	MIW	Naive
(c) Multivariate t	I	500	**2.78**	3.25	24.48	19.10
		1000	**1.39**	1.53	13.40	17.75
	LF	500	**12.65**	14.72	52.77	69.64
		1000	**6.71**	8.43	25.66	66.49
	AR	500	**20.54**	23.2	43.22	64.07
		1000	**13.53**	18.41	21.42	59.81
	EXP	500	**11.56**	14.12	37.68	43.57
		1000	**8.19**	11.97	18.22	41.66
(d) Multivariate Laplace	I	500	**1.81**	2.32	9.60	10.31
		1000	**0.97**	1.20	4.20	8.86
	LF	500	**7.33**	10.30	17.52	41.89
		1000	**3.99**	5.28	7.65	40.93
	AR	500	**9.79**	14.13	15.64	35.50
		1000	**5.54**	9.32	6.59	34.91
	EXP	500	**7.26**	9.90	13.93	23.71
		1000	**3.90**	5.12	5.19	22.78

Table S.2: Mean integrated squared error (MISE) performance of MLFA (mixtures of latent factor analyzers) and MIW (mixtures with inverse Wishart priors) density deconvolution models for **conditionally heteroscedastic** errors compared with a naive method that ignores measurement errors for different measurement error distributions. See Section 2 and Section S.6 for additional details. The minimum value in each row is highlighted.

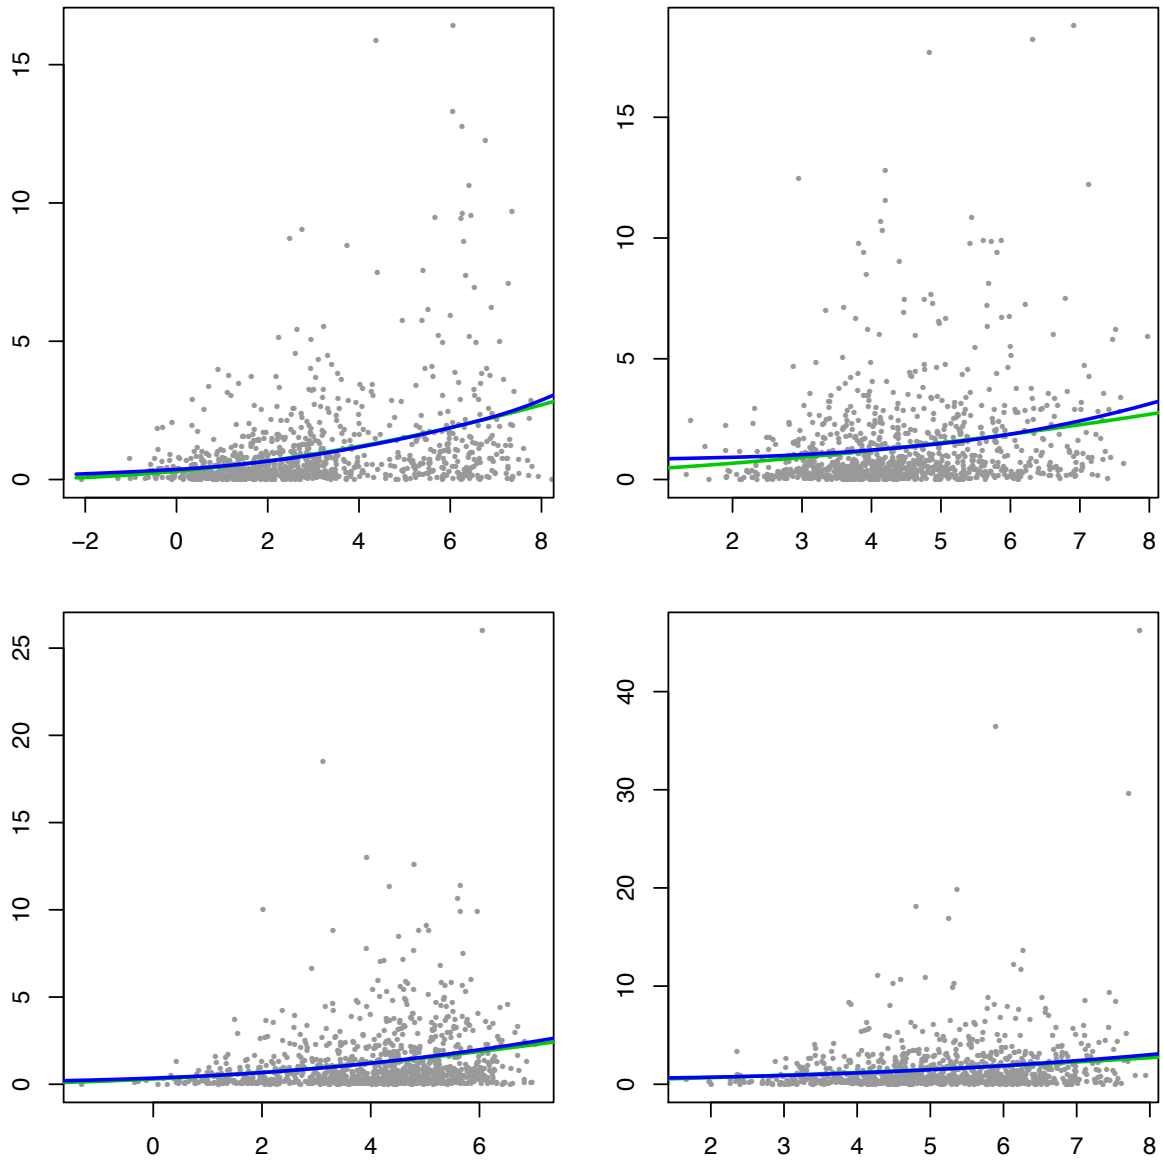

Figure S.13: Results for the variance functions $s^2(X)$ produced by the univariate density deconvolution method for each component of $\mathbf{X}$ for conditionally heteroscedastic multivariate Laplace ($f_{\boldsymbol{\epsilon}}^{(4)}$) distributed measurement errors with sample size $n = 1000$, $m_i = 3$ replicates for each subject and identity matrix (I) for the component specific covariance matrices. The results correspond to the data set that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets for the MIW (mixtures with inverse Wishart priors) method. For each component of $\mathbf{X}$, the true variance function is $s^2(X) = (1 + X/4)^2$. See Section 2.2.2 and Section 6 for additional details. In each panel, the true (lighter shaded lines) and the estimated (darker shaded lines) variance functions are superimposed over a plot of subject specific sample means vs subject specific sample variances. The figure is in color in the electronic version of this article.

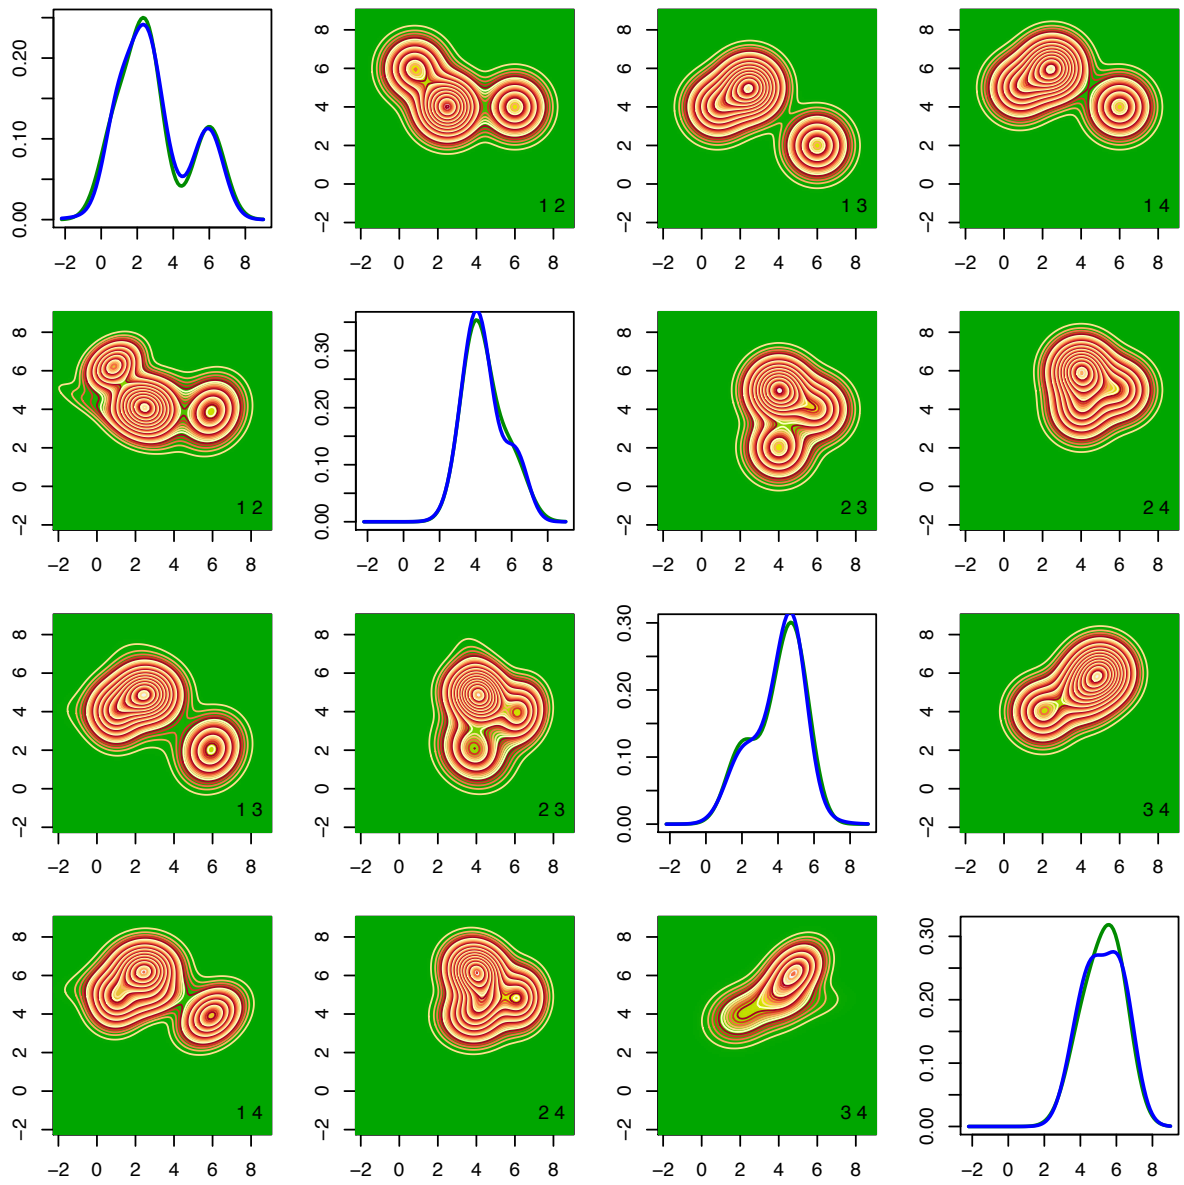

Figure S.14: Results for the $f_{\mathbf{X}}$ produced by the MIW (mixtures with inverse Wishart priors) method for conditionally heteroscedastic multivariate Laplace ($f_{\boldsymbol{\epsilon}}^{(4)}$) distributed measurement errors with sample size $n = 1000$, $m_i = 3$ replicates for each subject and identity matrix (I) for the component specific covariance matrices. The results correspond to the data set that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets. See Section 7 and Section S.6 for additional details. The upper triangular panels show the contour plots of the true two dimensional marginal densities. The lower triangular diagonally opposite panels show the corresponding estimates. The diagonal panels show the true (lighter shaded lines) and the estimated (darker shaded lines) one dimensional marginals. The figure is in color in the electronic version of this article.

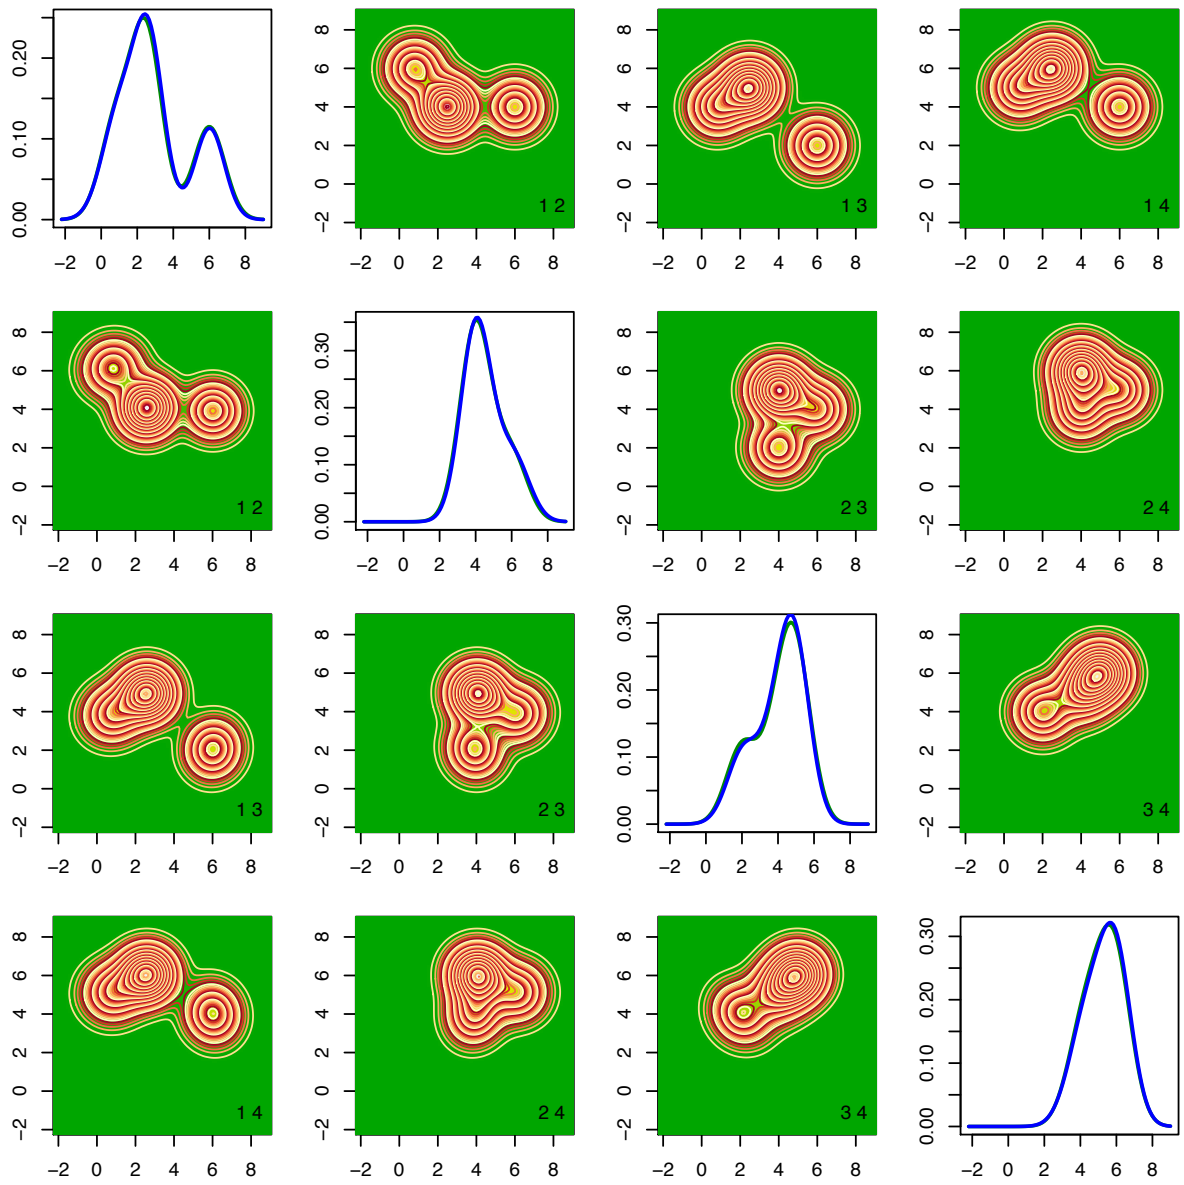

Figure S.15: Results for the $f_{\mathbf{X}}$ produced by the MLFAD (mixtures of latent factor analyzers) method for conditionally heteroscedastic multivariate Laplace ($f_{\boldsymbol{\epsilon}}^{(4)}$) distributed measurement errors with sample size $n = 1000$, $m_i = 3$ replicates for each subject and identity matrix (I) for the component specific covariance matrices. The results correspond to the data set that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets. See Section 7 and Section S.6 for additional details. The upper triangular panels show the contour plots of the true two dimensional marginal densities. The lower triangular diagonally opposite panels show the corresponding estimates. The diagonal panels show the true (lighter shaded lines) and the estimated (darker shaded lines) one dimensional marginals. The figure is in color in the electronic version of this article.

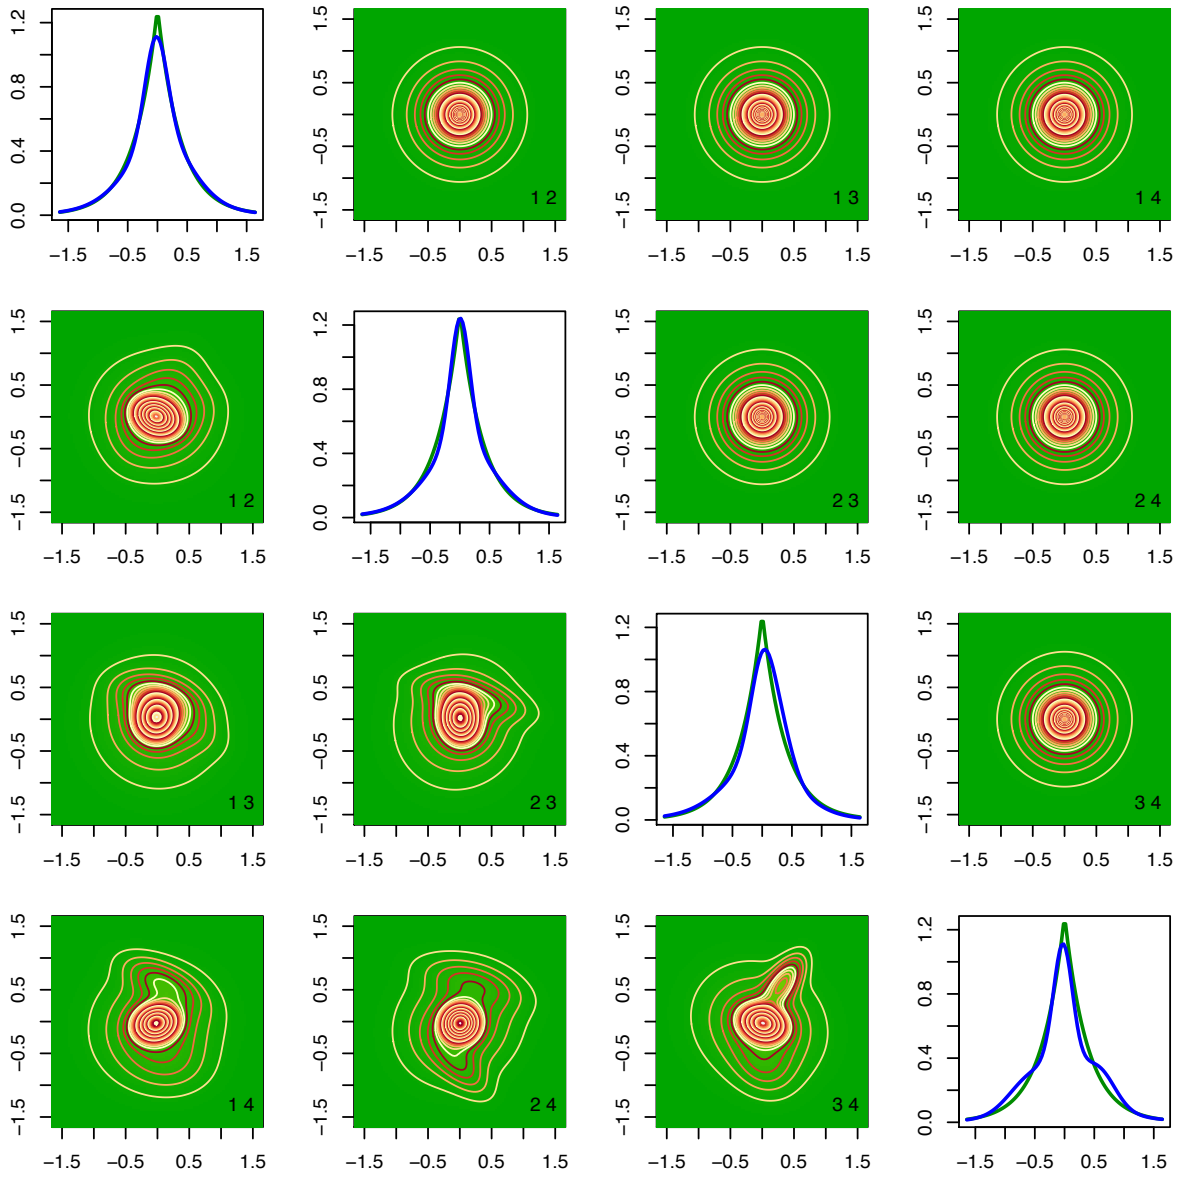

Figure S.16: Results for the density of the scaled errors $f_{\boldsymbol{\epsilon}}$ produced by the MIW (mixtures with inverse Wishart priors) method for conditionally heteroscedastic multivariate Laplace ($f_{\boldsymbol{\epsilon}}^{(4)}$) distributed measurement errors with sample size $n = 1000$, $m_i = 3$ replicates for each subject and identity matrix (I) for the component specific covariance matrices. The results correspond to the data set that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets. See Section 7 and Section S.6 for additional details. The upper triangular panels show the contour plots of the true two dimensional marginal densities. The lower triangular diagonally opposite panels show the corresponding estimates. The diagonal panels show the true (lighter shaded lines) and the estimated (darker shaded lines) one dimensional marginals. The figure is in color in the electronic version of this article.

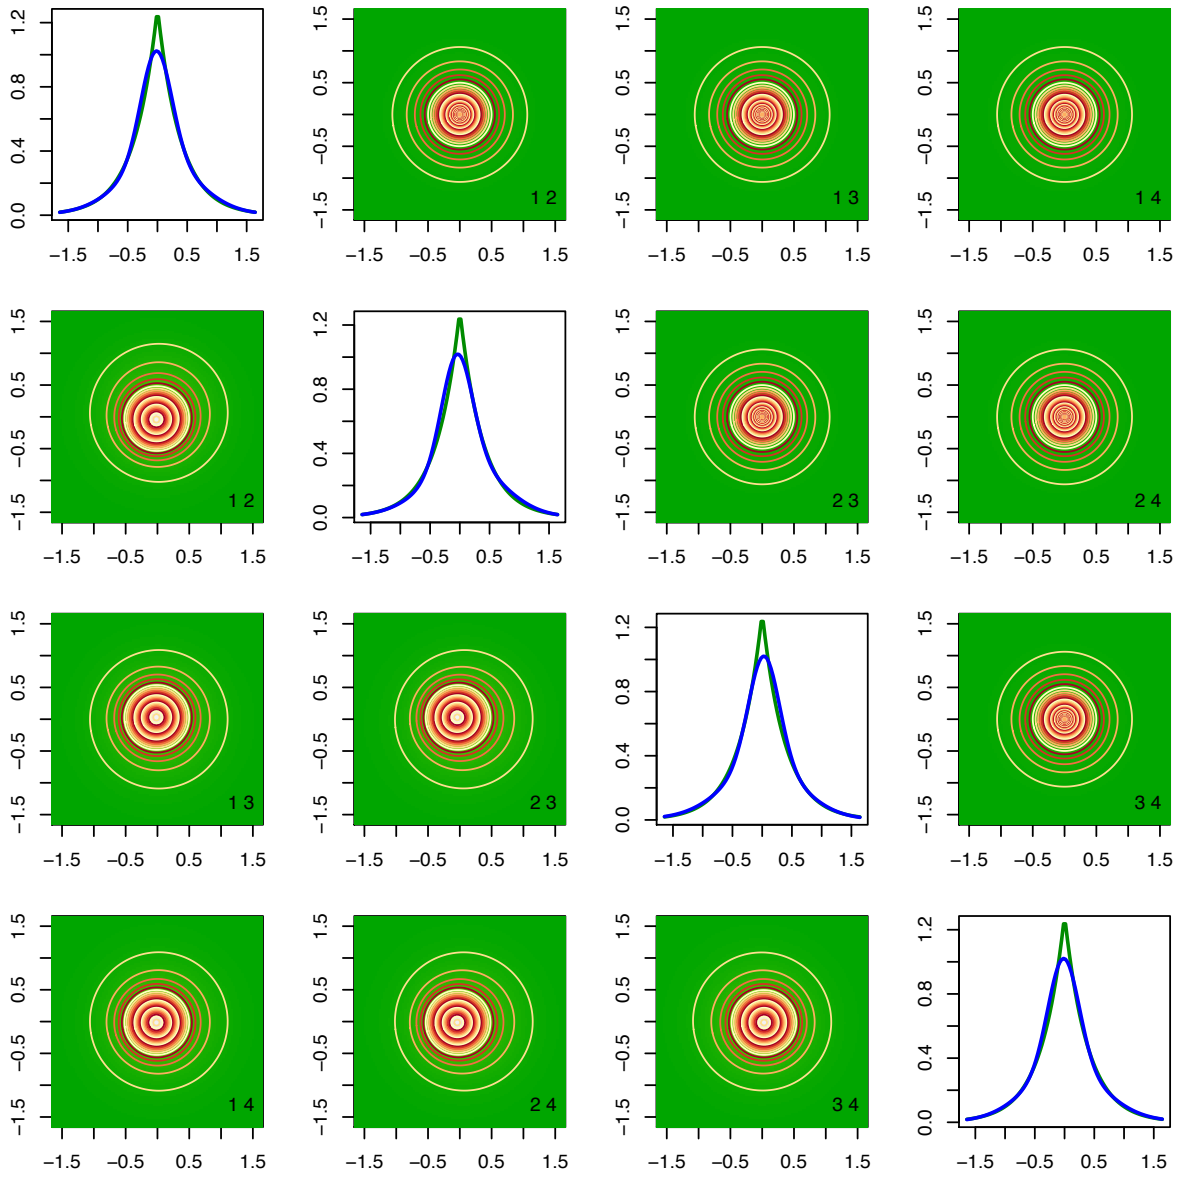

Figure S.17: Results for the density of the scaled errors $f_{\boldsymbol{\epsilon}}$ produced by the MLFAD (mixtures of latent factor analyzers) method for conditionally heteroscedastic multivariate Laplace ($f_{\boldsymbol{\epsilon}}^{(4)}$) distributed measurement errors with sample size $n = 1000$, $m_i = 3$ replicates for each subject and identity matrix (I) for the component specific covariance matrices. The results correspond to the data set that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets. See Section 7 and Section S.6 for additional details. The upper triangular panels show the contour plots of the true two dimensional marginal densities. The lower triangular diagonally opposite panels show the corresponding estimates. The diagonal panels show the true (lighter shaded lines) and the estimated (darker shaded lines) one dimensional marginals. The figure is in color in the electronic version of this article.

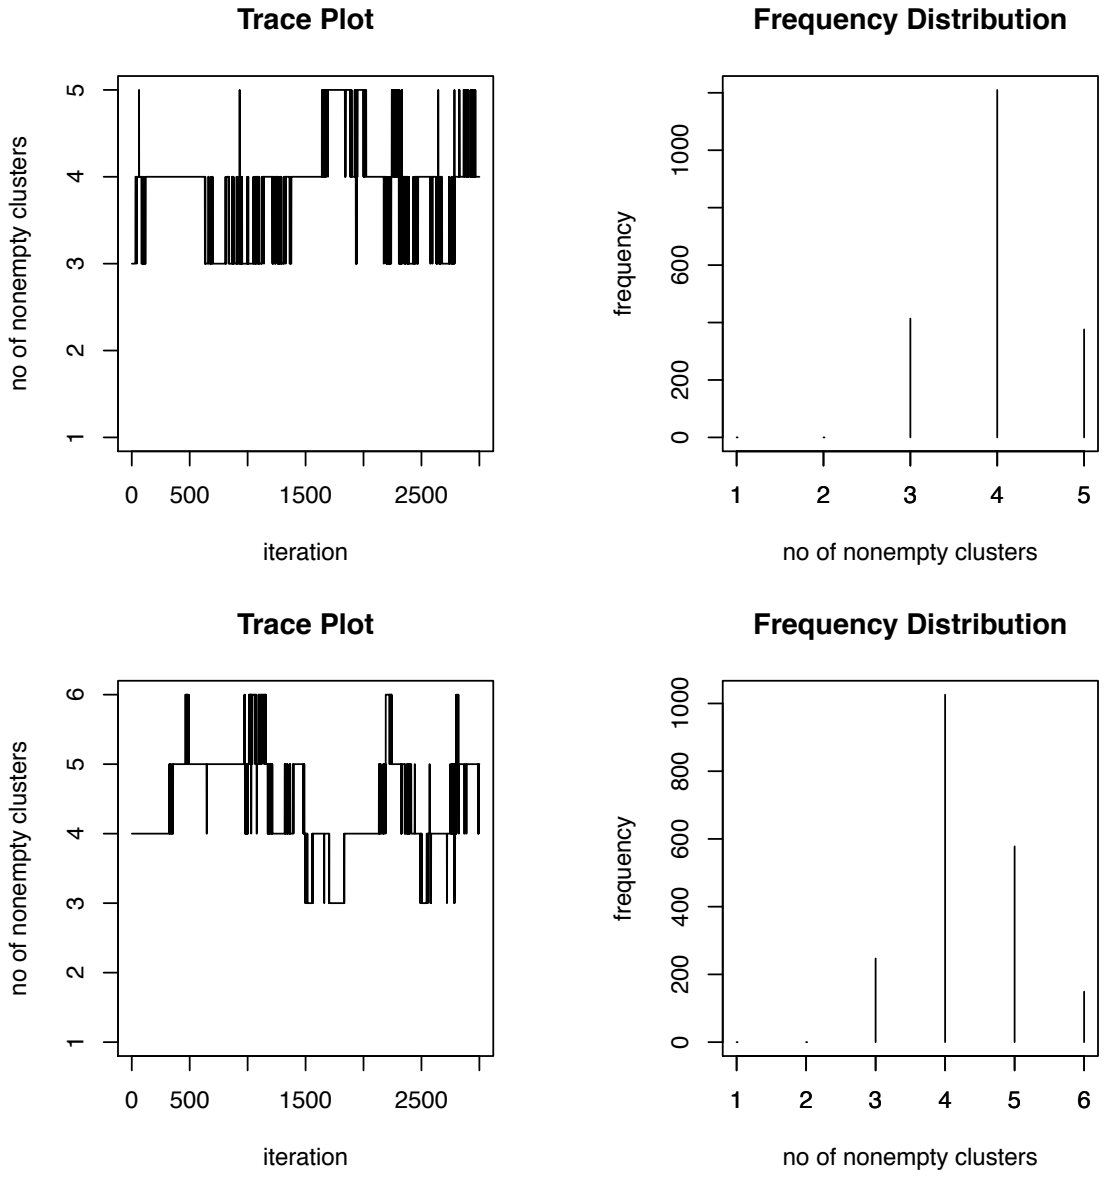

Figure S.18: Trace plots and frequency distributions of the number of nonempty clusters produced by the MIW (mixtures with inverse Wishart priors) method for conditionally heteroscedastic multivariate Laplace ($f_{\boldsymbol{\epsilon}}^{(4)}$) distributed measurement errors with sample size $n = 1000$, $m_i = 3$ replicates for each subject and identity matrix (I) for the component specific covariance matrices. See Section 7 and Section S.6 for additional details. The upper panels are for the $f_{\mathbf{X}}$ and the lower panels are for the density of the scaled errors $f_{\boldsymbol{\epsilon}}$. The results correspond to the simulation instance that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets, when the number of mixture components for both $f_{\mathbf{X}}$ and $f_{\boldsymbol{\epsilon}}$ were kept fixed at $K_{\mathbf{X}} = 5$ and $K_{\boldsymbol{\epsilon}} = 6$, respectively.

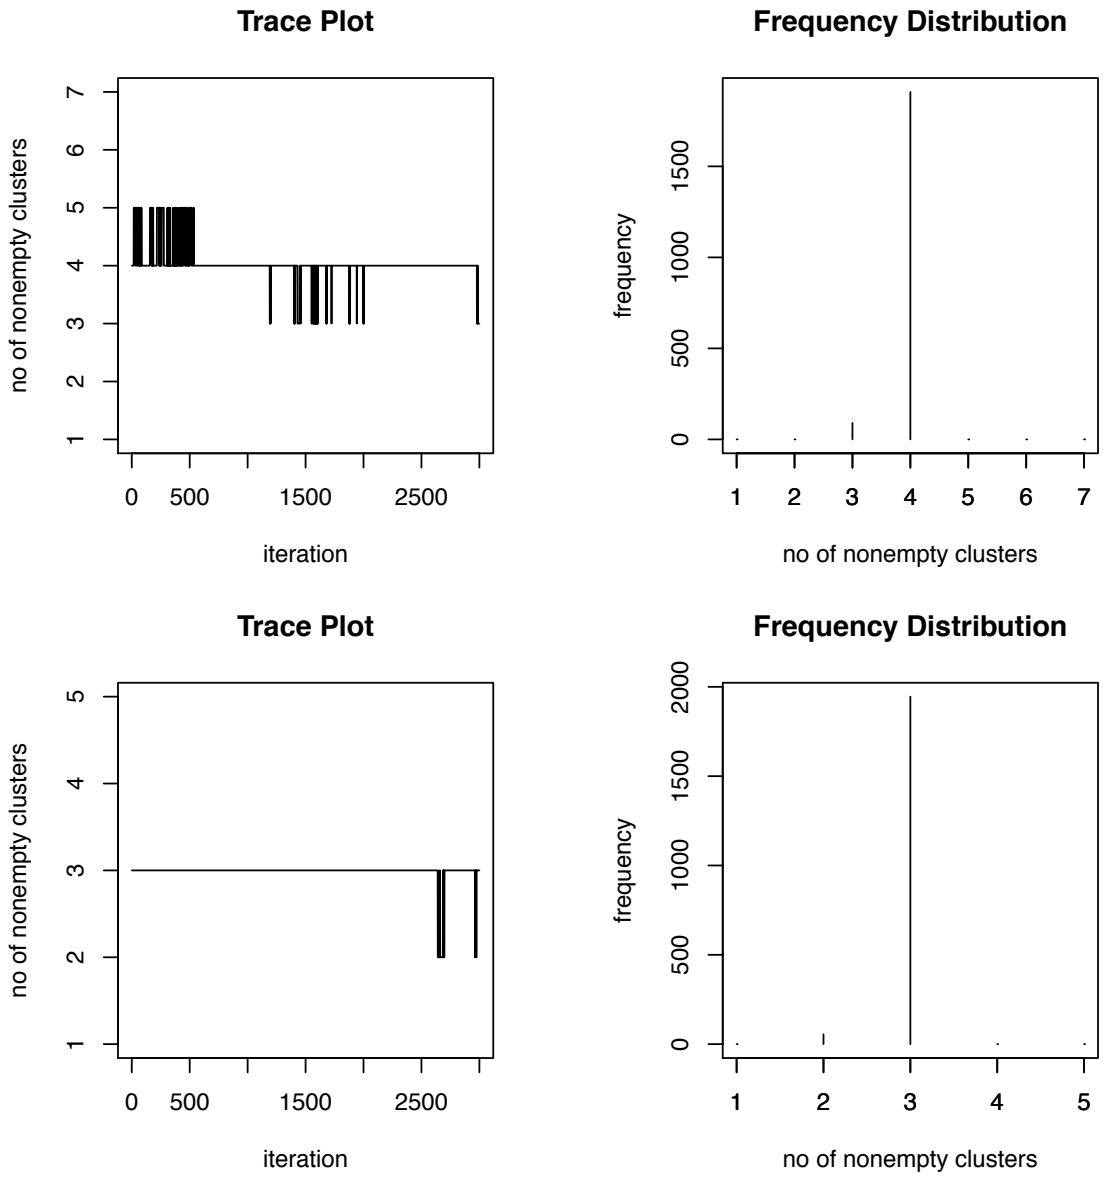

Figure S.19: Trace plots and frequency distributions of the number of nonempty clusters produced by the MLFAD (mixtures of latent factor analyzers) method for conditionally heteroscedastic multivariate Laplace ($f_{\boldsymbol{\epsilon}}^{(4)}$) distributed measurement errors with sample size $n = 1000$, $m_i = 3$ replicates for each subject and identity matrix (I) for the component specific covariance matrices. See Section 7 and Section S.6 for additional details. The upper panels are for the $f_{\mathbf{X}}$ and the lower panels are for the density of the scaled errors $f_{\boldsymbol{\epsilon}}$. The results correspond to the simulation instance that produced the median of the estimated integrated squared errors (ISE) out of a total of 100 simulated data sets, when the number of mixture components for $f_{\mathbf{X}}$ and $f_{\boldsymbol{\epsilon}}$ were kept fixed at $K_{\mathbf{X}} = 7$ and $K_{\boldsymbol{\epsilon}} = 5$, respectively.

Additional References

de Boor, C. (2000). *A Practical Guide to Splines*. New York: Springer.

Eltoft, T., Kim, T. and Lee, T. W. (2006). On the Multivariate Laplace Distribution. *IEEE Signal Processing Letters*, 13, 300-303.

Escobar, M. D. and West, M. (1995). Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association*, 90, 577-588.

Eubank, R. L. and Hart, J. D. (1992). Testing goodness-of-fit in regression via order selection criteria. *Annals of Statistics*, 20, 1412-1425.

Ferguson, T. F. (1973). A Bayesian analysis of some nonparametric problems. *Annals of Statistics*, 1, 209-230.

Ghosh, J. K. and Ramamoorthi, R. V. (2010). *Bayesian nonparametrics*. New York: *Springer*.

Hastie, D. I., Liverani, S. and Richrdson, S. (2013). Sampling from Dirichlet process mixture models with unknown concentration parameter: mixing issues in large data implementations. *Unpublished manuscript*.

Ishwaran, H. and Zarepour, M. (2000). Markov chain Monte Carlo in approximate Dirichlet and beta two-parameter process hierarchical models. *Biometrika*, 87, 371-390.

Ishwaran, H. and Zarepour, M. (2002). Exact and approximate sum-representations for the Dirichlet process. *Canadian Journal of Statistics*, 30, 269-283.

Kotz, S. and Nadarajah, S. (2004). Multivariate t Distributions and Their Applications. Cambridge: Cambridge University Press.

Neal, R. M. (2000). Markov chain sampling methods for Dirichlet process mixture models. *Journal of Computational and Graphical Statistics*, 9, 249-265.

Norets, A. and Pelenis, J. (2012). Bayesian modeling of joint and conditional distributions. *Journal of Econometrics*, 168, 332-346.

Pati, D. and Dunson, D. (2013). Bayesian nonparametric regression with varying residual density. *Annals of the Institute of Statistical Mathematics*, 66, 1-13.